arXiv:1612.08490v1 [stat.ME] 27 Dec 2016

Decorrelation of Covariates for High Dimensional Sparse Regression *

Jianqing Fan, Yuan Ke, Kaizheng Wang

Princeton University

December 28, 2016

Abstract

This paper studies the model selection consistency for high dimensional sparse regression with correlated covariates for a general class of regularized M estimators. The commonly-used model selection methods fail to consistently recover the true model when the covariates are not weakly correlated. This paper proposes a consistent model selection strategy named Factor Adjusted Decorrelation (FAD) for high dimensional sparse regression when the covariate dependence can be reduced through factor models. By separating the latent factors from idiosyncratic components, we transform the problem from model selection with highly correlated covariates to that with weakly correlated variables. We show that FAD can achieve model selection consistency as well as optimal rates of convergence under mild conditions. Numerical studies show FAD has nice finite sample performance in terms of both models selection and out-of-sample prediction. Moreover, FAD is a flexible method in a sense that it pays no price for weakly correlated and uncorrelated cases. The proposed method is applicable to a wide range of high dimensional sparse regression.

Key words: High dimension; Model selection consistency; Correlated covariates; Factor model; Regularized M estimator.

*Supported in part by NSF Grants DMS-1206464 and DMS-1406266 and R01-GM072611-9.

1 Introduction

High dimensional datasets are widely encountered in many areas such as genomics, medicine, machine learning, finance, marketing, among others. Model selection has become a popular and fundamental approach in high dimensional regression problems when the underlying model has a sparse representation. Parsimonious models are preferable due to their simplicity and interpretability. In addition, identifying significant covariates can improve the prediction accuracy. Therefore, over the past two decades, many model selection methods have been developed. A majority part of them are based on the regularized M-estimation approach including the Lasso (Tibshirani, 1996), the SCAD (Fan and Li, 2001), the elastic net (Zou and Hastie, 2005), and the Dantzig selector (Candes and Tao, 2007), among others. These methods have attracted a large amount of theoretical and algorithmic studies. See Donoho and Elad (2003), Fan and Peng (2004), Efron et al. (2004), Meinshausen and Bühlmann (2006), Zhao and Yu (2006), Fan and Lv (2008), Zou and Li (2008), Bickel et al. (2009), Wainwright (2009), Zhang (2010), and references therein.

In model selection studies, an issue of both theoretical and practical interests is to evaluate how well the sparse solution associates with the true model. This has been studied via an oracle property by Fan and Li (2001) for folded concave penalty functions. The model selection consistency of the convex L^1 regularization problem has been long studied due to its popularity and simplicity. Zhao and Yu (2006) characterized the model selection consistency of the Lasso by studying a stronger but technically more convenient sign consistency property and derived the *irrepresentable condition*. Meinshausen and Bühlmann (2006) showed a set of conditions under which the Lasso achieves model selection consistency for Gaussian graphic models. Bunea (2008) and Ravikumar et al. (2010) studied the model selection consistency for the regularized logistic regression. Van De Geer and Müller (2012) studied the high dimensional generalized linear model with L^1 regularization. They showed that the estimator has no false positive discoveries under a θ-*irrepresentable condition*. Lee et al. (2015) proposed a generalized *irrepresentable condition* for regularized M-estimators.

In this paper, we study the model selection consistency for high dimensional sparse regression. The *irrepresentable condition* indicates that the L^1-regularized regression fails to consistently select the true model when any covariate in the inactive set is highly correlated with the active set. However, this condition is restricted in real applications. First, the *irrepresentable condition* itself is hard to verify in practice as underlying active and inactive sets are unknown. Second, many datasets are generated in a correlated manner. Economics studies (e.g. Stock and Watson, 2002; Bai and Ng, 2002) show that there exist strong co-movements among a large pool of macroeconomic variables. A stylized feature of the stock return data is the stock returns are widely cross-sectional correlated. Similarly, co-expressions of genomic data can make the irrepresentable condition to fail. Third, spurious correlations in high dimensions can be so high that can easily fail the condition.

To illustrate how the correlations among covariates influence the model selection result, we

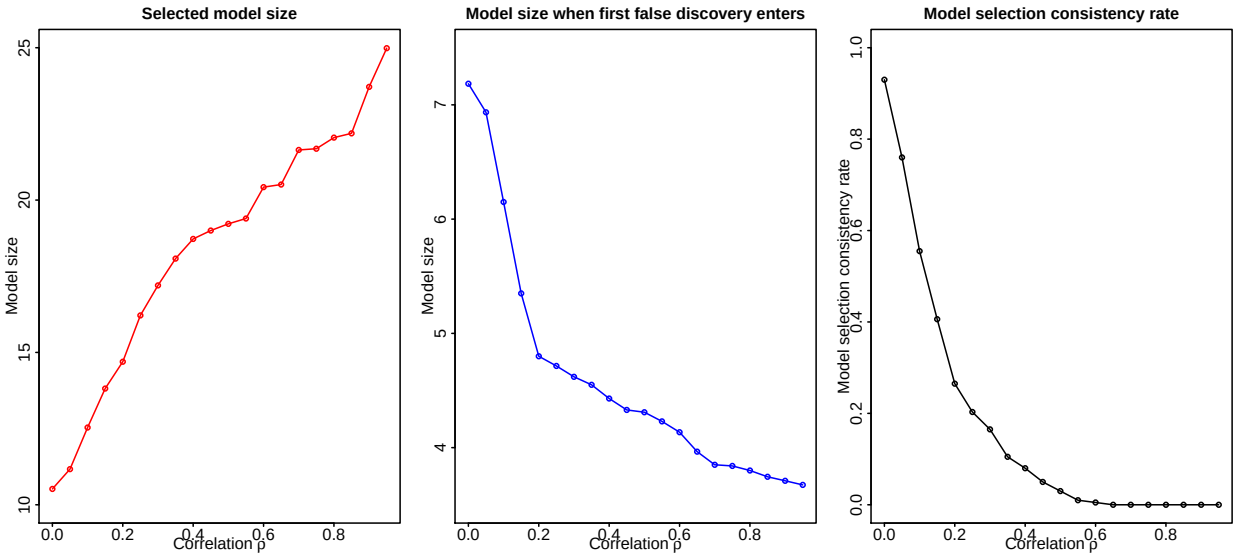

Figure 1: Lasso model selection results with respect to the correlations

consider Lasso with an equally correlated design. Take a sparse linear model $\mathbf{Y} = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\varepsilon}$ with the sample size $n = 100$, the dimensionality $p = 200$, $\boldsymbol{\beta}^* = (\beta_1, \cdots, \beta_{10}, \mathbf{0}^T_{(p-10)})^T$, and $\boldsymbol{\varepsilon} \sim N_n(\mathbf{0},\ 0.3\mathbf{I})$. The nonzero coefficients $\beta_1, \cdots, \beta_{10}$ are drawn from i.i.d. Uniform $[2, 5]$. The covariates $\mathbf{X} = (\mathbf{x}_1,\ \cdots,\ \mathbf{x}_p)^T$ are drawn from the normal distribution $N_p(\mathbf{0},\ \boldsymbol{\Sigma})$ where $\boldsymbol{\Sigma}$ is a correlation matrix with all off-diagonal elements ρ for some $\rho \in [0,\ 1)$. Let ρ increase from 0 to 0.95 by a step size 0.05. For each given ρ, we simulate 200 replications and calculate the average model size selected by Lasso, the average model size when the first false discovery ($\mathbf{x}_j,\ j > 10$) enters the solution path and the model selection consistency rate. As shown in Figure 1, the correlation influences the model selection results in the following three aspects: (i) selected model size, (ii) early selection of false variables, (iii) model selection consistency rates. Therefore, when the covariates are highly correlated, there is little hope to exactly recover the active set from the solution path of Lasso. As to be shown later, the correlation has similar adverse impacts on other model selection methods (e.g. SCAD and elastic net).

The influence of correlation in model selection is well recognized in existing literature (e.g. Lee et al., 2015; Su et al., 2017). Efforts have been made on how to reduce the impact of correlations for sparse linear models. Paul et al. (2008) studied the model selection when the response variable and covariates are connected via a low-dimensional latent variable model and proposed a "preconditioning" method. Jia and Rohe (2015) introduced the "Puffer transform" on the design matrix to comply with the *irrepresentable condition*. Kneip and Sarda (2011) and Wang (2012) studied the factor augmented sparse linear model.

To overcome the the aforementioned problems caused by the correlation, this paper proposes

a consistent model selection strategy named Factor Adjusted Decorrelation (FAD) for high dimensional sparse regression when the covariates can be decorrelated via a few pervasive latent factors. More precisely, let x_{ij} be the ith $(i = 1, \cdots, n)$ observation of the jth $(j = 1, \cdots, p)$ covariate, and assume that $\mathbf{x}_i = (x_{i1}, \cdots, x_{ip})^T$ follows an approximate factor model

$$\mathbf{x}_i = \mathbf{B}\mathbf{f}_i + \mathbf{u}_i, \tag{1.1}$$

where $\mathbf{f}_i$ is a $K \times 1$ vector of latent factors, $\mathbf{B}$ is a $p \times K$ matrix of loadings, and $\mathbf{u}_i$ is a $p \times 1$ vector of idiosyncratic components that are uncorrelated with $\mathbf{f}_i$. The strategy of FAD is to first estimate the the parameters in approximate factor model (1.1). Denote $\widehat{\mathbf{f}}_i$ and $\widehat{\mathbf{B}}$ the obtained estimators of the factors and loadings respectively. Then by identifying the highly correlated low rank part by $\widehat{\mathbf{B}}\widehat{\mathbf{f}}_i$, we transform the problem from model selection with highly correlated covariates in $\mathbf{x}_i$ to model selection with weakly correlated or uncorrelated idiosyncratic components $\widehat{\mathbf{u}}_i := \mathbf{x}_i - \widehat{\mathbf{B}}\widehat{\mathbf{f}}_i$. The second step amounts to solve a regularized profile likelihood problem. We study FAD in details by providing theoretical guarantees that FAD can achieve model selection consistency as well as estimation consistency under mild conditions. Moreover, both theoretical and numerical studies show FAD is a flexible method in a sense that it pays no price for weakly correlated cases. This property makes FAD very flexible when the underlying correlations between active and inactive covariates are unknown.

FAD is applicable to a wide range of high dimensional sparse regression related problems include but not limited to linear model, generalized linear model, Gaussian graphic model, robust linear model and group Lasso. For the sparse linear regression, the proposed approach is equivalent to project the response variable and covariates onto the linear space orthogonal to the one spanned by the estimated factors. Existing algorithms that yield solution paths of Lasso can be directly applied in the second step. To demonstrate the finite sample performance of FAD, we study two simulated and one empirical examples. The numerical results show FAD can consistently select the true model even when the covariates are highly correlated while existing methods like Lasso, SCAD and elastic net fail to do so.

Various methods have been studied to estimate the approximate factor model. Principal components analysis (PCA, Stock and Watson, 2002; Fan et al., 2013) is among one of the most popular ones. Data-driven estimation methods of the number of factors have been studied in extensive literature, such as Bai and Ng (2002), Hallin and Liška (2007), Lam and Yao (2012), and Ahn and Horenstein (2013) among others. Recently, a large amount of literature contributed to the asymptotic analysis of PCA under the ultra-high dimensional regime including Johnstone and Lu (2009), Shen et al. (2016) and Wang and Fan (2017), among others.

The rest of the paper is organized as follows. Section 2 overviews the problem setup including regularized M-estimator of sparse regression, the *irrepresentable condition* and approximate factor

models. Section 3 introduces the consistent model selection methodology of FAD and studies the sparse generalized linear model as a showcase example. Some issues related to the estimation of approximate factor models will be discussed in Section 3 as well. Section 4 presents the general theoretical results. Section 5 provides simulation studies and Section 6 analyzes an empirical dataset. The appendix contains the technical proofs.

Here are some notations that will be used throughout the paper. $\mathbf{I}_n$ denotes the $n \times n$ identity matrix; $\mathbf{0}_{n\times m}$ refers to the $n \times m$ zero matrix. For a matrix $\mathbf{M}$, we denote its matrix entry-wise max norm as $\|\mathbf{M}\|_{\max} = \max_{i,j}\{|M_{ij}|\}$ and denote by $\|\mathbf{M}\|_F$ and $\|\mathbf{M}\|_p$ its Frobenius and induced p-norms, respectively. For $\mathbf{M} \in \mathbb{R}^{n\times m}$, $I \subseteq [n]$ and $J \subseteq [m]$, define $\mathbf{M}_{IJ} = (\mathbf{M}_{ij})_{i\in I, j\in J}$, $\mathbf{M}_{I\cdot} = (\mathbf{M}_{ij})_{i\in I, j\in[m]}$ and $\mathbf{M}_{\cdot J} = (\mathbf{M}_{ij})_{i\in[n], j\in J}$. For a vector $\mathbf{v} \in \mathbb{R}^p$ and $S \subseteq [p]$, define $\mathbf{v}_S = (\mathbf{v}_i)_{i\in S}$ to be its subvector. Let ∇ and ∇^2 be the gradient and Hessian operators. For $f : \mathbb{R}^p \to \mathbb{R}$ and $I, J \in [p]$, define $\nabla_I f(x) = (\nabla f(x))_I$ and $\nabla^2_{IJ} f(x) = (\nabla^2 f(x))_{IJ}$.

2 Problem Setup

2.1 Regularized M-estimator

Let us consider a family of high dimensional sparse regression with the following settings. Suppose $\mathbf{Y} = (y_1, \cdots y_n)^T \in \mathbb{R}^n$ is a response vector of n independent observations with each y_i sampled from a probability distribution $\mathbb{P}(\theta_i)$ parametrized by θ_i and $\mathbf{X} = (\mathbf{x}_1, \cdots \mathbf{x}_n)^T \in \mathbb{R}^{n\times p}$ is a design matrix of covariates. Further we assume $\theta_i = \mathbf{x}_i^T\boldsymbol{\beta}^*$, where $\boldsymbol{\beta}^* \in \mathbb{R}^p$ is a sparse parameter vector with s non-zero elements. Let $L_n(\mathbf{Y}, \mathbf{X}\boldsymbol{\beta})$ be some convex and differentiable loss function that assigns a cost to any parameter $\boldsymbol{\beta} \in \mathbb{R}^p$ and $\boldsymbol{\beta}^*$ be the unique minimizer of the population risk $\mathbb{E}L_n(\mathbf{Y}, \mathbf{X}\boldsymbol{\beta})$. Under the high-dimensional regime, it is natural to estimate $\boldsymbol{\beta}^*$ via a regularized M-estimator as follows:

$$\widehat{\boldsymbol{\beta}} \in \operatorname*{argmin}_{\boldsymbol{\beta}\in\mathbb{R}^p} \{L_n(\mathbf{Y}, \mathbf{X}\boldsymbol{\beta}) + \lambda R_n(\boldsymbol{\beta})\}, \tag{2.1}$$

where $R_n : \mathbb{R}^p \to \mathbb{R}_+$ is a norm that penalizes the use of a nonsparse vector $\boldsymbol{\beta}$ and $\lambda > 0$ is a tuning parameter.

A special case of this problem is the L^1 penalized likelihood estimation of generalized linear models. Suppose the conditional density function of Y given $X = \mathbf{x}$ is a member of the exponential family, i.e.

$$f(y|\mathbf{x}, \boldsymbol{\beta}^*) \propto \exp\{y\mathbf{x}^T\boldsymbol{\beta}^* - b(\mathbf{x}^T\boldsymbol{\beta}^*) + c(y)\},$$

where $b(\cdot)$ and $c(\cdot)$ are known functions and $\boldsymbol{\beta}^*$ is an unknown sparse coefficient vector of interest. Taking the loss function to be the negative log-likelihood function and the penality function to be

the L^1 norm, the regularized M-estimator of $\boldsymbol{\beta}^*$ admits the form

$$\widehat{\boldsymbol{\beta}} \in \min_{\boldsymbol{\beta}\in\mathbb{R}^p} \left\{ \frac{1}{n}\sum_{i=1}^{n}[-y_i\mathbf{x}_i^T\boldsymbol{\beta} + b(\mathbf{x}_i^T\boldsymbol{\beta})] + \lambda\|\boldsymbol{\beta}\|_1 \right\}. \tag{2.2}$$

2.2 Irrepresentable condition

We expect a good estimator of (2.1) to achieve estimation consistency as well as selection consistency. The former one requires $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\| \xrightarrow{p} \mathbf{0}$ as $n \to \infty$; while the latter one requires $P\{\text{supp}(\widehat{\boldsymbol{\beta}}) = \text{supp}(\boldsymbol{\beta}^*)\} \to 1$ as $n \to \infty$. In general, the estimation consistency does not imply the selection consistency and vice versa. To study the selection consistency, we consider a stronger condition named general sign consistency as follows.

Definition 2.1 (Sign consistency)**.** *An estimate $\widehat{\boldsymbol{\beta}}$ is sign consistent with $\boldsymbol{\beta}^*$ if $\exists\lambda \geq 0$ such that*

$$\lim_{n\to\infty} P(\text{sign}(\widehat{\boldsymbol{\beta}}) = \text{sign}(\boldsymbol{\beta}^*)) = 1.$$

Zhao and Yu (2006) studied the Lasso estimator and showed there exists an *irrepresentable condition* which is sufficient and almost necessary for both sign and estimation consistencies. Without loss of generality, we assume $\text{supp}(\boldsymbol{\beta}^*) = [s] = S$. Denote $\mathbf{X}_S$ and $\mathbf{X}_{S^c}$ as the submatrices of $\mathbf{X}$ defined by its first s columns and the rest $p - s$ columns, respectively. Then the *irrepresentable condition* requires some $\tau \in (0, 1)$, such that

$$\|\mathbf{X}_{S^c}^T\mathbf{X}_S(\mathbf{X}_S^T\mathbf{X}_S)^{-1}\|_\infty < \tau. \tag{2.3}$$

For general regularized M-estimator (2.1) to achieve both sign and estimation consistencies, Lee et al. (2015) proposed a generalized *irrepresentable condition*. When applied to the L^1 regularizer, it becomes

$$\|\nabla^2_{S^cS}L(\boldsymbol{\beta}^*)(\nabla^2_{SS}L(\boldsymbol{\beta}^*))^{-1}\|_\infty < \tau, \tag{2.4}$$

for some $\tau \in (0, 1)$. It is easy to check (2.4) is equivalent to (2.3) under the Lasso case. The generalized *irrepresentable condition* will easily get violated when there exists strong correlations between active and inactive variables. Even if it holds, the key parameter τ can be very close to zero, making it hard to select the correct model and obtain small estimation errors at the same time.

2.3 Approximate factor model

To go beyond the weakly correlated assumption, a natural extension is conditional weak correlation (Fan and Lv, 2008; Fan et al., 2013). Suppose covariates are dependent through latent common

factors. Given these common factors, the idiosyncratic components are weakly correlated. We assume $\mathbf{x}_i \in \mathbb{R}^p$ follows the approximate factor model (1.1). Notice $\mathbf{x}_i$ is the only observable in (1.1). Throughout the paper, we assume K is independent of n.

Denote $\mathbf{F} = (\mathbf{f}_1, \cdots, \mathbf{f}_n)^T \in \mathbb{R}^{n\times K}$ and $\mathbf{U} = (\mathbf{u}_1, \cdots, \mathbf{u}_n)^T \in \mathbb{R}^{n\times p}$. Then (1.1) can be written in a more compact matrix form:

$$\mathbf{X} = \mathbf{F}\mathbf{B}^T + \mathbf{U}. \tag{2.5}$$

Let $\mathbf{\Sigma}$ be the covariance matrix of $\{\mathbf{x}_i\}_{i=1}^n$ and $\lambda_1, \cdots, \lambda_p$ be the eigenvalues of $\mathbf{\Sigma}$ in the descending order. In order to consistently identify the dependence part induced by the common factors from the sparse part that corresponds to the weakly correlated idiosyncratic components, we consider the following spiked covariance model for $\mathbf{\Sigma}$ (Wang and Fan, 2017).

Assumption 2.1. *Assume $p > n$ and $\lambda_1 > \lambda_2 > \cdots > \lambda_K > \lambda_{K+1} \geq \cdots \geq \lambda_p > 0$.*
*(i) For the **spiked** part $1 \leq j \leq K$, $c_j = p/(n\lambda_j)$ is bounded and the spiked eigenvalues are well separated, i.e. $\exists \delta_0 > 0$ such that $\min_{j\leq K}(\lambda_j - \lambda_{j+1})/\lambda_j \geq \delta_0$.*
*(ii) For the **non-spiked** part $j > K$, the eigenvalues are uniformly bounded, i.e. there exists constants $c_0, C_0 > 0$ such that $c_0 \leq \lambda_j \leq C_0$. In addition $(p-K)^{-1}\sum_{j=k+1}^p \lambda_j = \bar{c} + o(n^{-1/2})$, where $\bar{c}$ is a bounded constant.*

The Assumption 2.1 allows $p/n \to \infty$ in any manner and the spiked eigenvalues $\{\lambda_j\}_{j=1}^K$ are allowed to grow slower than p so long as $c_j := p/(n\lambda_j)$ is bounded. This assumption is weaker than the pervasiveness assumption $\lambda_j = O(p)$ required in existing literature (e.g. Stock and Watson, 2002; Bai and Ng, 2002; Fan et al., 2013).

3 Model selection with FAD

3.1 Methodology

By the approximate factor model (2.5),

$$\mathbf{X}\boldsymbol{\beta} = \mathbf{F}\mathbf{B}^T\boldsymbol{\beta} + \mathbf{U}\boldsymbol{\beta} := \mathbf{F}\boldsymbol{\gamma} + \mathbf{U}\boldsymbol{\beta},$$

where $\boldsymbol{\gamma} = \mathbf{B}^T\boldsymbol{\beta} \in \mathbb{R}^K$. By regarding $\boldsymbol{\gamma}$ as nuisance parameters, the regularized M-estimator (2.1) can be rewritten as

$$\widehat{\boldsymbol{\beta}} \in \underset{\boldsymbol{\gamma}\in\mathbb{R}^K,\ \boldsymbol{\beta}\in\mathbb{R}^p}{\operatorname{argmin}} \left\{L_n(\mathbf{Y}, \mathbf{F}\boldsymbol{\gamma} + \mathbf{U}\boldsymbol{\beta}) + \lambda R_n(\boldsymbol{\beta})\right\}. \tag{3.1}$$

By construction, $\mathbf{U}$ has now much weaker correlation than $\mathbf{X}$. This penalized profile likelihood removes the effect of strong correlations caused by the latent factors. It can be implemented as follows:

Step 1: Initial estimation. Fit the approximate factor model (2.5) and denote $\widehat{\mathbf{B}}$, $\widehat{\mathbf{F}}$ and $\widehat{\mathbf{U}} = \mathbf{X} - \widehat{\mathbf{F}}\widehat{\mathbf{B}}^T$ the obtained estimators of $\mathbf{B}$, $\mathbf{F}$ and $\mathbf{U}$ respectively.

Step 2: Augmented M-estimation. Define $\widehat{\mathbf{W}} = (\widehat{\mathbf{F}},\ \widehat{\mathbf{U}})$ and $\boldsymbol{\theta} = (\boldsymbol{\gamma}^T,\ \boldsymbol{\beta}^T)^T$. Then $\widehat{\boldsymbol{\beta}}$ can be obtained by solving the following augmented problem

$$\widehat{\boldsymbol{\theta}} \in \underset{\boldsymbol{\theta}\in\mathbb{R}^{K+p}}{\operatorname{argmin}} \left\{ L_n(\mathbf{Y}, \widehat{\mathbf{W}}\boldsymbol{\theta}) + \lambda R_n(\boldsymbol{\theta}_{[K^c]}) \right\}. \tag{3.2}$$

Suppose $\mathbf{U}$ is independent of $\mathbf{F}$ and the columns of $\mathbf{U}$ are weakly correlated, then the columns in $\mathbf{W} = (\mathbf{F}, \mathbf{U})$ are weakly correlated covariates as long as $\mathbf{F}$ and $\mathbf{U}$ are well estimated. Hence, we successfully transform the problem from model selection with highly correlated covariates $\mathbf{X}$ in (2.1) to model selection with weakly correlated or uncorrelated covariates $\mathbf{W}$. The augmented problem (3.2) is a convex optimization problem which can be minimized via many existing convex optimization algorithms, for example coordinate descent(e.g. Friedman et al., 2010) and ADMM.

3.2 Example: sparse linear model

Now we illustrate the FAD procedure using sparse linear regression. In this case, we have

$$\mathbf{Y} = \mathbf{X}\boldsymbol{\beta}^* + \boldsymbol{\varepsilon} = \widehat{\mathbf{F}}\widehat{\mathbf{B}}^T\boldsymbol{\beta}^* + \widehat{\mathbf{U}}\boldsymbol{\beta}^* + \boldsymbol{\varepsilon}, \tag{3.3}$$

The augmented M-estimator (3.2) for the sparse linear model is of the following form:

$$\widehat{\boldsymbol{\beta}} \in \underset{\boldsymbol{\beta}\in\mathbb{R}^p, \boldsymbol{\gamma}\in\mathbb{R}^K}{\operatorname{argmin}} \left\{ \frac{1}{2n}\|\mathbf{Y} - \widehat{\mathbf{F}}\boldsymbol{\gamma} - \widehat{\mathbf{U}}\boldsymbol{\beta}\|_2^2 + \lambda\|\boldsymbol{\beta}\|_1 \right\}.$$

Solving the least-squares problem with respect to $\boldsymbol{\gamma}$, we have the penalized profile least-squares solution

$$\widehat{\boldsymbol{\beta}} \in \underset{\boldsymbol{\beta}\in\mathbb{R}^p}{\operatorname{argmin}} \left\{ \frac{1}{2n}\|(\mathbf{I}_n - \widehat{\mathbf{P}})(\mathbf{Y} - \widehat{\mathbf{U}}\boldsymbol{\beta})\|_2^2 + \lambda\|\boldsymbol{\beta}\|_1 \right\}, \tag{3.4}$$

where $\widehat{\mathbf{P}} = \widehat{\mathbf{F}}(\widehat{\mathbf{F}}^T\widehat{\mathbf{F}})^{-1}\widehat{\mathbf{F}}^T$ is the $n \times n$ projection matrix onto the column space of $\widehat{\mathbf{F}}$. As the decorrelation step does not depend on the choice of the regularizer $R(\cdot)$, FAD can be applied to a wide range of penalized least squares problems such as SCAD, group Lasso, elastic net, fused Lasso, and so on.

A special case of our FAD method is the removal of intercept in regression. To do this, we center the covariates in $\mathbf{X}$ and add an all-one column to the design matrix. We can view this as letting $\widehat{\mathbf{F}} = \mathbf{1}_{n\times 1}$, $\widehat{\mathbf{B}} = \frac{1}{n}\mathbf{X}^T\widehat{\mathbf{F}}$, $\mathbf{U} = \mathbf{X} - \widehat{\mathbf{F}}\widehat{\mathbf{B}}^T$. The Lasso is implemented on the centered covariates and responses.

There is another way to understand this method. By left multiplying the projection matrix

$(\mathbf{I}_n - \widehat{\mathbf{P}})$ to both sides of (3.3), we have

$$(\mathbf{I}_n - \widehat{\mathbf{P}})\mathbf{Y} = (\mathbf{I}_n - \widehat{\mathbf{P}})\widehat{\mathbf{U}}\boldsymbol{\beta}^* + (\mathbf{I}_n - \widehat{\mathbf{P}})\boldsymbol{\varepsilon}, \tag{3.5}$$

where $(\mathbf{I}_n - \widehat{\mathbf{P}})\widehat{\mathbf{U}}$ can be treated as the decorrelated design matrix and $(\mathbf{I}_n - \widehat{\mathbf{P}})\mathbf{Y}$ is the corresponding response variable. From (3.5) we see that the method in Kneip and Sarda (2011) coincides with FAD in the linear case. However, the projection-based representation only makes sense in sparse linear regression. In contrast, our idea of profile likelihood directly generalizes to more general problems.

3.3 Estimating factor models

Principal component analysis (PCA) is frequently used to estimate latent factors for model (2.5). The estimated matrix of latent factors $\widehat{\mathbf{F}}$ is $\sqrt{n}$ times the eigenvectors corresponding to the K largest eigenvalues of the $n \times n$ matrix $\mathbf{X}\mathbf{X}^T$. Using the normalization $\mathbf{F}^T\mathbf{F}/n = \mathbf{I}_K$ yields $\widehat{\mathbf{B}} = \mathbf{X}^T\widehat{\mathbf{F}}/n$.

Now we introduce the asymptotic properties of estimated factors and idiosyncratic components. The following Lemma is similar to the results developed in the Appendix C of Wang and Fan (2017) and hence we omit the proof.

Assumption 3.1. $\|\boldsymbol{\Lambda}_K^{-1/2}\mathbf{B}^T\mathbf{B}\boldsymbol{\Lambda}_K^{-1/2} - \boldsymbol{\Omega}_0\|_2 = o(1)$ *for some* Ω_0 *with eigenvalues bounded from above and blow, where* $\boldsymbol{\Lambda}_K = diag(\lambda_1, \cdots, \lambda_K)$.

Assumption 3.2. *(i)* $\{\mathbf{f}_i, \mathbf{u}_i\}_{i=1}^n$ *are independent and identically distributed with* $\mathrm{E}(\mathbf{u}_i) = \mathbf{0}$ *and* $\mathrm{E}(\mathbf{f}_i) = \mathbf{0}$.
(ii) There exists positive constants c_1 *and* c_2 *such that* $\lambda_{\min}(\boldsymbol{\Sigma}_u) > c_1$, $\|\Sigma_u\|_\infty < c_2$.
(iii) There exist positive constants r_1, r_2, b_1 *and* b_2 *such that for* $\delta > 0$, $j \le p$, $k \le K$,

$$\mathrm{P}(|u_{ij}| > \delta) \le exp(-(\delta/b_1)^{r_1}) \quad and \quad \mathrm{P}(|f_{ik}| > \delta) \le exp(-(\delta/b_2)^{r_2}).$$

(iv) There exists $M > 0$ *such that for all* $j \le p$, $k \le K$, $|b_j k| \le M\sqrt{\lambda_j/p}$.
(v) $\sqrt{p}(\log n)^{1/r_2} = o(\lambda_m)$.

Lemma 3.1. *Under Assumptions 2.1, 3.1 and 3.2, there exists nonsingular* $\mathbf{H}_0 \in \mathbb{R}^{K\times K}$ *such that:*
(i) $\|\widehat{\mathbf{F}}\mathbf{H}_0 - \mathbf{F}\|_{\max} = O_p((\frac{1}{\sqrt{\lambda_K}} + \frac{p}{n\lambda_K} + \frac{\sqrt{p}}{\lambda_K})(\log n)^{2/r_2})$;
(ii) $\|\widehat{\mathbf{F}}\mathbf{H}_0 - \mathbf{F}\|_F = O_p(\frac{p}{\sqrt{n}\lambda_K} + \sqrt{\frac{n}{\lambda_K}})$;
(iii) $\max_{1\le j\le p} n^{-1}\sum_{i=1}^n |\widehat{\mathbf{U}}_{ji} - \mathbf{U}_{ji}|^2 = O_p(\frac{\log p}{n} + \frac{1}{p})$; *and*
(iv) $\|\widehat{\mathbf{U}} - \mathbf{U}\|_{\max} = o_p(1)$.

A practical issue arises on how to choose the number of factors. We adapt the ratio method for the numerical studies in this paper. Let $\lambda_k(\mathbf{X}\mathbf{X}^T)$ be the kth largest eigenvalue of $\mathbf{X}\mathbf{X}^T$ and K_{max}

be a prescribed upper bound. The number of factors can be consistently estimated by (Lam and Yao, 2012; Ahn and Horenstein, 2013)

$$\widehat{K} = \underset{k \le K_{max}}{\operatorname{argmax}} \frac{\lambda_k(\mathbf{X}\mathbf{X}^T)}{\lambda_{k+1}(\mathbf{X}\mathbf{X}^T)}. \tag{3.6}$$

Other viable method includes the information criteria in Bai and Ng (2002).

3.4 Decorrelated variable screening

Screening methods (e.g. Fan and Lv, 2008; Fan and Song, 2009; Wang and Leng, 2016) are computationally attractive and thus popular for ultra-high dimensional data analysis. However, the screening methods will recruit too many variables when there exist strong correlations among covariates (Fan and Lv, 2008; Wang and Leng, 2016). As an extension of FAD, we propose the following conditional variable screening method to tackle this problem.

Step 1: *Initial estimation.* We fit the approximate factor model (2.5) to obtain $\widehat{\mathbf{B}}$, $\widehat{\mathbf{F}}$ and $\widehat{\mathbf{U}} = \mathbf{X} - \widehat{\mathbf{F}}\widehat{\mathbf{B}}^T$.

Step 2: *Augmented marginal regression.* For $j \in [p]$, let $\widehat{\mathbf{v}}_j$ be the j-th column of $\mathbf{U}$ and

$$\widehat{\theta}_j = \underset{\boldsymbol{\gamma} \in \mathbb{R}^K, \theta \in \mathbb{R}}{\operatorname{argmin}} \; L_n(\mathbf{Y}, \widehat{\mathbf{F}}\boldsymbol{\gamma} + \widehat{\mathbf{v}}_j^T \theta). \tag{3.7}$$

Step 3 *Screening.* Sort the $\{\widehat{\theta}_j\}_{j=1}^p$ in terms of their absolute values, and take the largest ones.

For the sparse linear regression case, our screening method reduces to the factor-profiled screening proposed by Wang (2012).

4 Theoretical results

4.1 General results

We first present the general model selection results for the FAD estimator (3.2). Without loss of generality, we assume the first K variables are not penalized. Let $L_n : \mathbb{R}^{p+K} \to \mathbb{R}$ be a convex loss function, $\boldsymbol{\theta}^* \in \mathbb{R}^{p+K}$ and $\boldsymbol{\beta}^* = \boldsymbol{\theta}^*_{[K]^c}$ be the sparse sub-vector of interest. Then $\boldsymbol{\theta}^*$ and $\boldsymbol{\beta}^*$ are estimated via

$$\widehat{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta} \in \mathbb{R}^{K+p}}{\operatorname{argmin}} \{L_n(\boldsymbol{\theta}) + \lambda \|\boldsymbol{\theta}_{[K]^c}\|_1\} \quad \text{and} \quad \widehat{\boldsymbol{\beta}} = \widehat{\boldsymbol{\theta}}_{[K]^c},$$

respectively. Further, denote $S = \operatorname{supp}(\boldsymbol{\theta}^*)$, $S_1 = \operatorname{supp}(\boldsymbol{\beta}^*)$ and $S_2 = [p+K] \backslash S$.

Assumption 4.1. *(Smoothness)* $L_n(\boldsymbol{\theta}) \in C^2$ *and there exist* $A > 0$, $M > 0$ *such that* $\|\nabla^2_{\cdot S} L(\boldsymbol{\theta}) - \nabla^2_{\cdot S} L_n(\boldsymbol{\theta}^*)\|_\infty \leq M\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2$ *whenever* $\mathrm{supp}(\boldsymbol{\theta}) \subseteq S$ *and* $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| \leq A$.

Assumption 4.2. *(Restricted strong convexity) There exist* $\kappa_2 > \kappa_\infty > 0$ *such that* $\|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_\infty \leq \frac{1}{2\kappa_\infty}$ *and* $\|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_2 \leq \frac{1}{2\kappa_2}$.

Assumption 4.3. *(Irrepresentable condition)* $\|\nabla^2_{S_2 S} L_n(\boldsymbol{\theta}^*)(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_\infty \leq 1 - \tau$ *for some* $\tau \in (0, 1)$.

Assumptions 4.1 – 4.3 are standard in the studies of high-dimensional regularized estimators (e.g. Negahban et al., 2012; Lee et al., 2015). Based on them, we introduce the following theorem of L^p $(p = 1, 2, \infty)$ error bounds and sign consistency for the FAD estimator.

Theorem 4.1. *(i)* ***Error bounds*** *: Under Assumptions 4.1 – 4.3, if*

$$\frac{7}{\tau}\|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty < \lambda < \frac{\kappa_2}{4\sqrt{|S|}} \min\left\{A, \frac{\kappa_\infty \tau}{3M}\right\}, \tag{4.1}$$

then $\mathrm{supp}(\widehat{\boldsymbol{\theta}}) \subseteq S$ *and*

$$\begin{aligned}
&\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \leq \frac{3}{5\kappa_\infty}(\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \lambda), \\
&\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \leq \frac{2}{\kappa_2}(\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_2 + \lambda\sqrt{|S_1|}), \\
&\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \leq \min\left\{\frac{3}{5\kappa_\infty}(\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_1 + \lambda|S_1|), \frac{2\sqrt{|S|}}{\kappa_2}(\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_2 + \lambda\sqrt{|S_1|})\right\}.
\end{aligned}$$

(ii) ***Sign consistency*** *: In addition, if the following two conditions*

$$\begin{aligned}
&\min\{|\boldsymbol{\beta}^*_j| : \boldsymbol{\beta}^*_j \neq 0,\ j \in [p]\} > \frac{C}{\kappa_\infty \tau}\|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty, \\
&\|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty < \frac{\kappa_2 \tau}{7C\sqrt{|S|}} \min\left\{A, \frac{\kappa_\infty \tau}{3M}\right\}
\end{aligned} \tag{4.2}$$

hold for some $C \geq 5$, *then by taking* $\lambda \in (\frac{7}{\tau}\|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty, \frac{1}{\tau}(\frac{5C}{3} - 1)\|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty)$, *the estimator achieves the sign consistency* $\mathrm{sign}(\widehat{\boldsymbol{\beta}}) = \mathrm{sign}(\boldsymbol{\beta}^*)$.

Remark 4.1. Theorem 4.1 shows how the correlated covariates affect the sign consistency as well as error bounds. To achieve the sign consistency, the tuning parameter λ should scale with τ^{-1}. Therefore, the L^∞ and L^2 errors will scale with $(\kappa_\infty \tau)^{-1}$ and $(\kappa_2 \tau)^{-1}$, respectively. When the covariates are highly correlated, the *irrepresentable condition* will get violated or the parameter $\tau \in (0, 1)$ is very small. As a result, the model selection procedures will fail to achieve the sign

consistency and the error bounds will be suboptimal. On the other hand, the optimal error bounds require a small λ, which typically leads to an overfitted model. One can see, there is a trade-off between model selection and parameter estimation due to the existence of dependency.

Remark 4.2. The L^1 and L^2 error bounds in Theorem 4.1 depend on $|S_1|$, the number of active variables. They stem from the bias induced by the penalty. To reduce the bias, it is desirable to penalize as few active variables as possible. This phenomena motivates FAD to adapt a penalized profile likelihood form by not imposing penalty on the nuisance parameter $\boldsymbol{\gamma}$.

As discussed in Remark 4.1, when the covariates are highly correlated, the *irrepresentable condition* may not hold, or has a very small τ. This makes the model selection consistency either very hard to achieve or incompatible with low estimation error bounds. Therefore, the FAD strategy can improve the model selection consistency and reduce the estimation error bounds if $\mathbf{X}$ can be decomposed into $\mathbf{F}\mathbf{B}^T + \mathbf{U}$ such that $\mathbf{W} = (\mathbf{F}, \mathbf{U})$ is well-behaved. This is due to the fact that the *irrepresentable condition* is easier to hold with positive τ bounded away from zero after the decorrelation step. To this end, any effective decorrelation procedure can be incorporated into this frame work.

4.2 FAD with approximate factor model

Now we study the FAD estimator when the covariates $\mathbf{X}$ admits the approximate factor model (2.5). The oracle estimator that knows the true augmented covariates $\mathbf{W} = (\mathbf{F}, \mathbf{U})$ and consider the minimizer of

$$\min_{\boldsymbol{\theta}} \{L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}) + \lambda \|\boldsymbol{\theta}_{[K]^c}\|_1\}. \tag{4.3}$$

However, $\mathbf{W}$ is not observable and hence one need to use its estimator $\widehat{\mathbf{W}}$ and solve

$$\min_{\boldsymbol{\theta}} \{L_n(\mathbf{Y}, \widehat{\mathbf{W}}\boldsymbol{\theta}) + \lambda \|\boldsymbol{\theta}_{[K]^c}\|_1\}. \tag{4.4}$$

Below the error induced by the factor estimation will be studied carefully. To deliver a clear discussion on the assumptions and results, we focus on the FAD estimator for the generalized linear model (2.2).

Assumption 4.4. *$b(z) \in C^3(\mathbb{R}^2)$. For some constants M_2 and M_3, we have $0 \le b''(z) \le M_2$ and $|b'''(z)| \le M_3$, $\forall z$.*

Assumption 4.5. *There exist $\kappa_2 > \kappa_\infty > 0$ and $\tau \in (0, 1)$ such that*

$$\begin{aligned} &\|(\nabla^2_{SS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\ell \le \frac{1}{4\kappa_\ell}, \quad \text{for } \ell = 2 \text{ and } \infty. \\ &\|\nabla^2_{S_2 S} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*)(\nabla^2_{SS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \le 1 - 2\tau. \end{aligned} \tag{4.5}$$

Assumption 4.6. *(Estimation of factor model)* $\|\mathbf{W}\|_{\max} \le \frac{M_0}{2}$ *for some* $M_0 > 0$. *Besides, there exist* $K \times K$ *nonsingular matrix* $\mathbf{H}_0$, *and* $\mathbf{H} = \begin{pmatrix} \mathbf{H}_0 & \mathbf{0}_{K\times p} \\ \mathbf{0}_{p\times K} & \mathbf{I}_p \end{pmatrix}$ *such that for* $\overline{\mathbf{W}} = \widehat{\mathbf{W}}\mathbf{H}$, *we have* $\|\overline{\mathbf{W}} - \mathbf{W}\|_{\max} \le \frac{M_0}{2}$ *and*

$$\max_{1\le j\le p+K} \Big(\frac{1}{n}\sum_{i=1}^{n} |\overline{\mathbf{W}}_{ij} - \mathbf{W}_{ij}|^2\Big)^{1/2} \le \frac{2\kappa_\infty \tau}{3M_0 M_2 |S|}. \tag{4.6}$$

Remark 4.3. (i) Assumption 4.4 holds for a large family of generalized linear models. For example, linear model with $b(z) = \frac{1}{2}z^2$, $M_2 = 1$ and $M_3 = 0$; logistic modelwith $b(z) = \log(1 + e^z)$ and finite M_2, M_3. (ii) Note that the first inequality in (4.5) involves only a small matrix and holds easily, and the second inequality there is related to the generalized irrespresentable condition. Standard concentration inequalities yield that Assumption 4.5 holds with high probability if $\mathbb{E}[\nabla^2 L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*)]$ satisfies similar conditions. (iii)Under Assumptions 2.1, 3.1, 3.2 and $\lambda_K \gg \max\Big\{n, \frac{p}{\sqrt{n\log p}}, \frac{p(\log n)^{2/r_2}}{n}, \sqrt{p}(\log n)^{2/r_2}\Big\}$, Lemma 3.1 implies the existence of $\mathbf{H}_0$ in Assumption 4.6 such that $\delta_0 = o_P(1)$ and

$$\max_{1\le j\le p+K} \Big(\frac{1}{n}\sum_{i=1}^{n} |(\widehat{\mathbf{W}}\mathbf{H})_{ij} - \mathbf{W}_{ij}|^2\Big)^{1/2} = O_P(\sqrt{\frac{\log p}{n}}).$$

Hence $|S|^2\frac{\log p}{n} = O(1)$ can guarantee Assumption 4.6 to hold with high probability.

Theorem 4.2. *Suppose Assumptions 4.4, 4.5 and 4.6 hold. Define* $M = M_0^3 M_3 |S|^{3/2}$ *and*

$$\varepsilon = \max_{1\le j\le K+p} \Big|\frac{1}{n}\sum_{i=1}^{n} \overline{w}_{ij}[-y_i + b'(\mathbf{x}_i^T\boldsymbol{\beta}^*)]\Big|. \tag{4.7}$$

If $\frac{7\varepsilon}{\tau} < \lambda < \frac{\kappa_2\kappa_\infty\tau}{12M\sqrt{|S|}}$, *then we have* $\mathrm{supp}(\widehat{\boldsymbol{\beta}}) \subseteq \mathrm{supp}(\boldsymbol{\beta}^*)$ *and*

$$\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_\infty \le \frac{6\lambda}{5\kappa_\infty}, \qquad \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2 \le \frac{4\lambda\sqrt{|S|}}{\kappa_2}, \qquad \|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_1 \le \frac{6\lambda|S|}{5\kappa_\infty}.$$

In addition, if $\varepsilon < \frac{\kappa_2\kappa_\infty\tau^2}{12CM\sqrt{|S|}}$ *and*

$$\min\{|\boldsymbol{\beta}_j^*| : \boldsymbol{\beta}_j^* \ne 0,\ j \in [p]\} > \frac{6C\varepsilon}{5\kappa_\infty\tau} \tag{4.8}$$

hold for some $C > 7$, *then by taking* $\lambda \in (\frac{7}{\tau}\varepsilon, \frac{C}{\tau}\varepsilon)$ *we can achieve the sign consistency* $\mathrm{sign}(\widehat{\boldsymbol{\beta}}) = \mathrm{sign}(\boldsymbol{\beta}^*)$.

Remark 4.4. Taking $\lambda \asymp \varepsilon$ one can achieve the sign consistency and $\|\widehat{\boldsymbol{\beta}}-\boldsymbol{\beta}^*\|_\infty = O(\varepsilon)$, $\|\widehat{\boldsymbol{\beta}}-\boldsymbol{\beta}^*\|_2 = O(\varepsilon\sqrt{|S|})$ and $\|\widehat{\boldsymbol{\beta}}-\boldsymbol{\beta}^*\|_1 = O(\varepsilon|S|)$. If $\varepsilon = O_P(\sqrt{\frac{\log p}{n}})$, then all the error bounds match the optimal ones in the literature (e.g. Bickel et al., 2009). Note that $\mathbb{E}y_i = b'(\mathbf{x}_i^T\boldsymbol{\beta}^*)$. Hence

$$\varepsilon = \max_{1\le j\le K+p} \Big|\frac{1}{n}\sum_{i=1}^{n} \overline{w}_{ij}[y_i - \mathbb{E}y_i]\Big| \tag{4.9}$$

is the supremum of an empirical process defined by $\{y_i\}_{i=1}^n$, whose size can be well controlled. When the covariates $\{\overline{w}_{ij}\}_{i\in[n],j\in[p+K]}$ are all bounded (Assumption 4.6) and $\{y_i\}_{i=1}^n$ are subgaussian, $\varepsilon = O_P(\sqrt{\frac{\log p}{n}})$.

5 Simulation study

5.1 Example 1: Linear regression

We study a simulated example for high dimensional sparse linear regression with correlated covariates. The correlation structure is calibrated from S&P 500 monthly excess returns between 1980 and 2012. Throughout the numerical studies of this paper, the tuning parameter λ is selected by the 10-fold cross validation. The model selection performance is measured by the model selection consistency rate and the sure screening rate. The former is the proportion of simulations that the selected model is identical to the true one and the latter is the proportion of simulations that the selected model contains all important variables.

Calibration and data generating process

We calculate the monthly excess returns for the stocks in S&P 500 index that have complete record from January 1980 to December 2012. The data, collected from CRSP[1] , contains 202 stocks with a time span of 396 months. We fit the data with a factor model and set the number of common factors $K = 3$. Denote $\tilde{\mathbf{\Lambda}}$, $\tilde{\mathbf{f}}_t$ and $\tilde{\mathbf{u}}_t$ the estimated loadings, factors and residuals, and let $\tilde{\mathbf{H}}$ be a $K \times K$ rotation matrix such that $(\tilde{\mathbf{\Lambda}}\tilde{\mathbf{H}})^T(\tilde{\mathbf{\Lambda}}\tilde{\mathbf{H}})/p = \mathbf{I}_K$. Then we calculate $\tilde{\boldsymbol{\mu}}_f$ (the sample mean of $\tilde{\mathbf{f}}_t$) and the sample covariance matrix $\mathrm{cov}(\tilde{\mathbf{H}}^{-1}\tilde{\mathbf{f}}_t)$. Furthermore, we calculate $\mathrm{cov}(\tilde{\mathbf{u}}_t)$ and let $\tilde{\sigma}_u^2$ be the average of the diagonal entries of $\mathrm{cov}(\tilde{\mathbf{u}}_t)$. The calibrated parameters are presented in Table 1.

We generate the covariates $\mathbf{X}$ from the factor model (2.5) with $K = 3$. We draw $\mathbf{\Lambda}$ from the i.i.d. standard Normal distribution, $\mathbf{f}_t$ from the trivariate Normal distribution $N_3(\tilde{\boldsymbol{\mu}}_f,\ \mathrm{cov}(\tilde{\mathbf{H}}^{-1}\tilde{f}_t))$, and $\mathbf{u}$ from the i.i.d. Normal distribution $N(0,\ \tilde{\sigma}_u^2)$. Then we generate $\mathbf{Y}$ from a sparse linear model with $\boldsymbol{\varepsilon}$ drawn from i.i.d. Normal distribution $N(0,\ 0.3)$. The true coefficients are $\boldsymbol{\beta}^* = (\beta_1, \cdots, \beta_{10}, \mathbf{0}_{(p-10)}^T)^T$, and the nonzero coefficients $\beta_1, \cdots, \beta_{10}$ are drawn from i.i.d. Uniform $[2, 5]$.

[1]Center for Research in Security Prices Database, see http://www.crsp.com/ for more details.

Table 1: Parameters calibrated from S&P 500 returns

$\tilde{\boldsymbol{\mu}}_f$	$\text{cov}(\tilde{\mathbf{H}}^{-1}\tilde{f}_t)$			$\tilde{\sigma}_u^2$
-0.878	0.230	0.234	0.050	0.110
0.265	0.234	0.934	0.015	
0.057	0.050	0.015	1.017	

Impacts of Irrepresentable Condition

First, we show Lasso performs poorly in terms of model selection consistency rate when the *irrepresentable condition* is violated while FAD can consistently select the correct model. Let $n = 100$ and $p = 500$. Denote by $\Gamma_\infty = \|\mathbf{X}_{S^c}^T\mathbf{X}_S(\mathbf{X}_S^T\mathbf{X}_S)^{-1}\|_\infty$. When $\Gamma_\infty < 1$ the *irrepresentable condition* holds otherwise it is violated. We simulate 10,000 replications. For each replication, we calculate Γ_∞ and apply both Lasso and FAD to do model selection. Then we calculate the model selection consistency rate within each small interval around Γ_∞ (a nonparametric smoothing). The results are presented in Figure 2. According to Figure 2, both FAD and Lasso have high model selection consistency rate when $\Gamma_\infty < 1$. This shows FAD does not pay any price under the weak correlation scenario. As Γ_∞ grows beyond 1, the correct model selection rate of Lasso drops quickly. When the *irrepresentable condition* is strongly violated (e.g. $\Gamma_\infty > 1.5$), the correct model selection rate of Lasso is close to zero. On the contrary, FAD has high selection consistency rates regardless of Γ_∞.

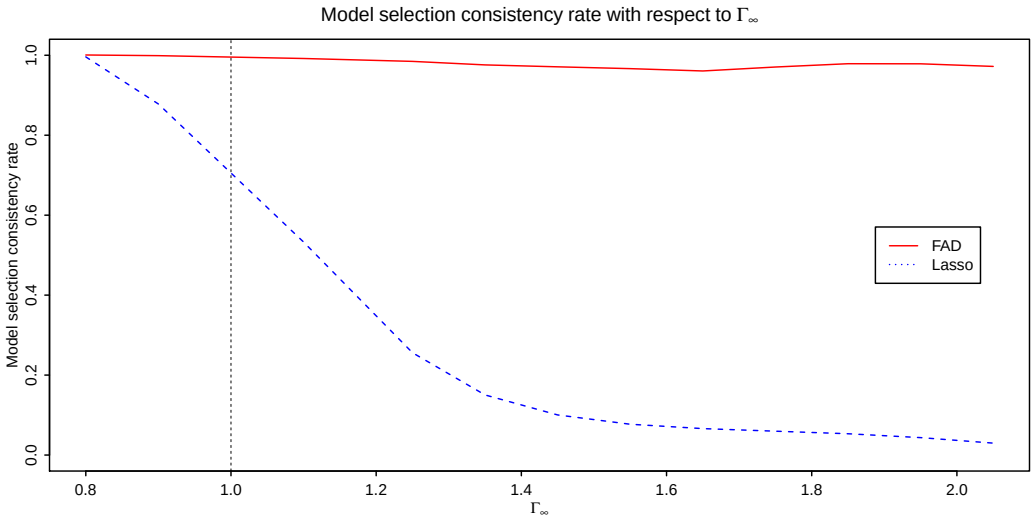

Figure 2: Relationship between model selection consistency rate and *irrepresentable condition*. Among the 10,000 replications, more than 9,500 replications have $\Gamma_\infty > 1$ and more than 8,000 replications have $\Gamma_\infty > 1.5$.

Impacts of sample size

Second, we examine the model selection consistency with a fixed dimensionality and an increasing sample size. We fix $p = 500$ and let n increase from 50 to 150. For each given sample size, we simulate 200 replications and calculate the model selection consistency rate and the sure screening rate for Lasso, SCAD, elastic net and FAD, respectively. For the elastic net, we set $\lambda_1 = \lambda_2 \equiv \lambda$. The results are presented in Figure 3 (a) and Figure 3 (b). Figure 3 (a) shows model selection consistency rates of Lasso, SCAD and elastic net do not converge to 1 fast when the sample size increases. The model selection consistency rate of FAD equals to one for all sample sizes greater than 100. So is the sure screening rate of FAD. The sure screening rates of the other three methods converge to one when the sample size increases. To demonstrate the prediction performance, we report the mean estimation error $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2$ for each method, which is a good indicator of the prediction error. The estimation errors are reported in Figure 3 (c). One can see, when the sample size is small, Lasso, SCAD and elastic net have large estimation errors due to they tend to select overfitted models.

Impacts of dimensionality

Third, we assess the model selection performance when the dimensionality p is growing beyond n and diverging. We fix $n = 100$ and let p grow from 200 to 1000. For each given p, we simulate 200 replications and calculate the model selection consistency rate of Lasso, SCAD, elastic net and FAD respectively. The model selection consistency rates are presented in Figure 4(a). According to Figure 4(a), the model selection consistency rate of FAD stays close to 1 even as p increases, whereas the model selection consistency rates for the other three methods drops quickly. Again, we report the estimation error or each model selection method in 4(b). The estimation error of FAD increases with dimensionality much slower than the other three methods.

5.2 Example 2: Logistic regression

We consider the following logistic regression model whose conditional probability function is:

$$\mathbb{P}[y_i = 1|\mathbf{X}_i] = \frac{\exp(\mathbf{X}_i^T\boldsymbol{\beta})}{1 + \exp(\mathbf{X}_i^T\boldsymbol{\beta})}, \quad i = 1, \cdots, N. \tag{5.1}$$

We set sample size $n = 200$, dimensionality $p = 300,\ 400,\ 500$. The true coefficients are set to be $\boldsymbol{\beta}^* = (\boldsymbol{\beta}_{(1)}^T,\ \mathbf{0})^T$ with $\boldsymbol{\beta}_{(1)} = (6,\ 5,\ 4)^T$. Hence the true model size is 3.

The covariates $\mathbf{X}$ are generated from one of the following three models:

(1) Factor model with $K = 3$. The factors, loading and idiosyncratic errors are generated from the i.i.d. standard Normal distribution.

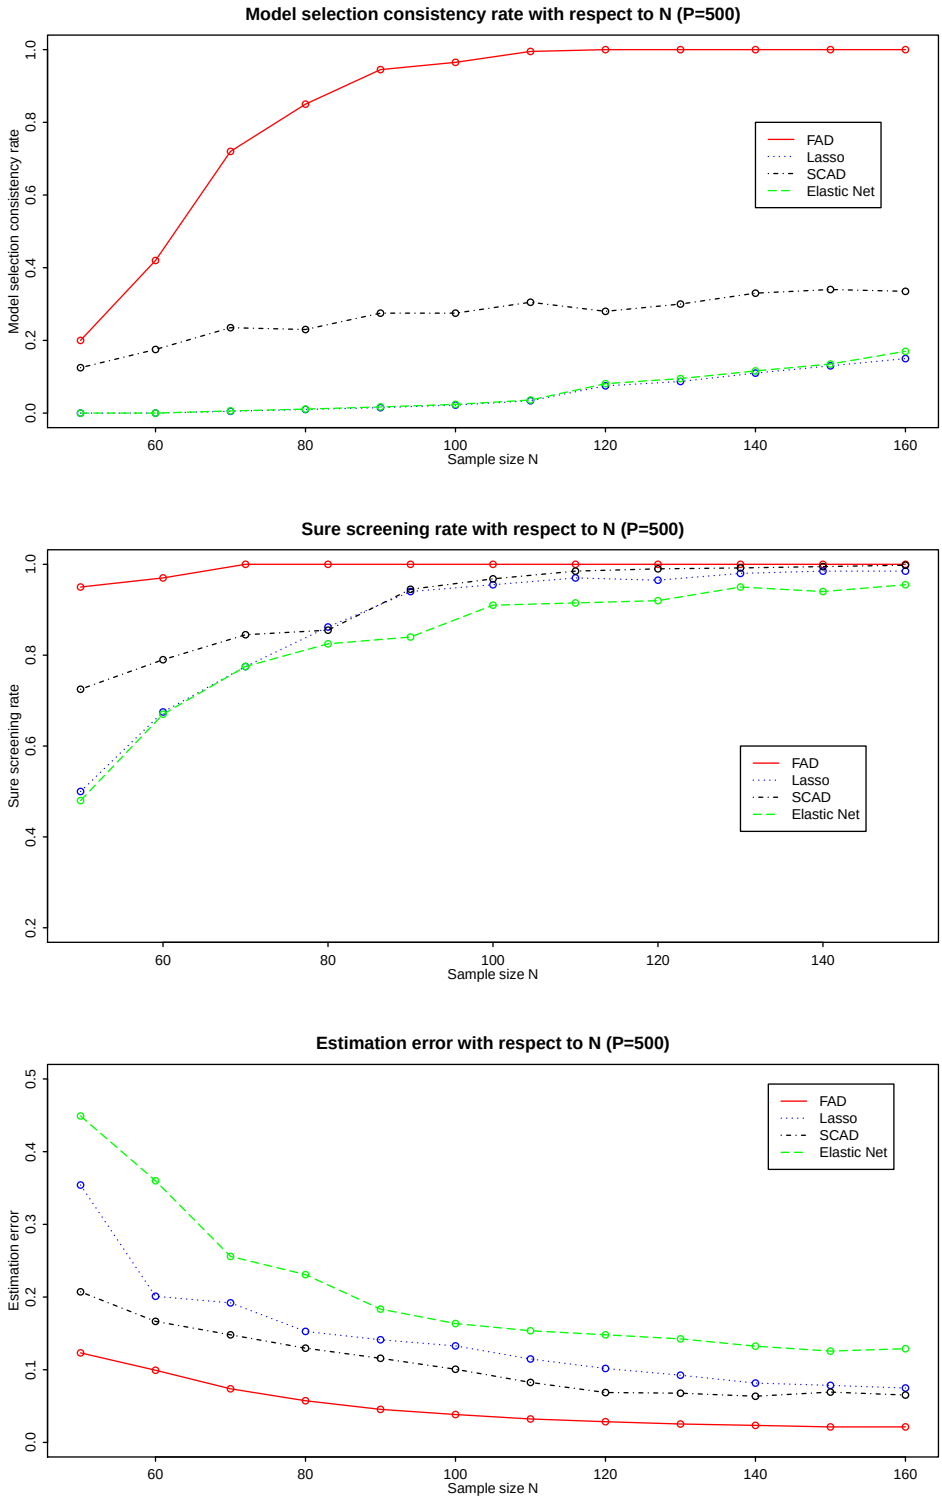

Figure 3: From above to blow: (a) Model selection consistency rate with fixed p and increasing n; (b) Sure screening rate with fixed p and increasing n; (c) Mean estimation error $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2$ with fixed p and increasing n.

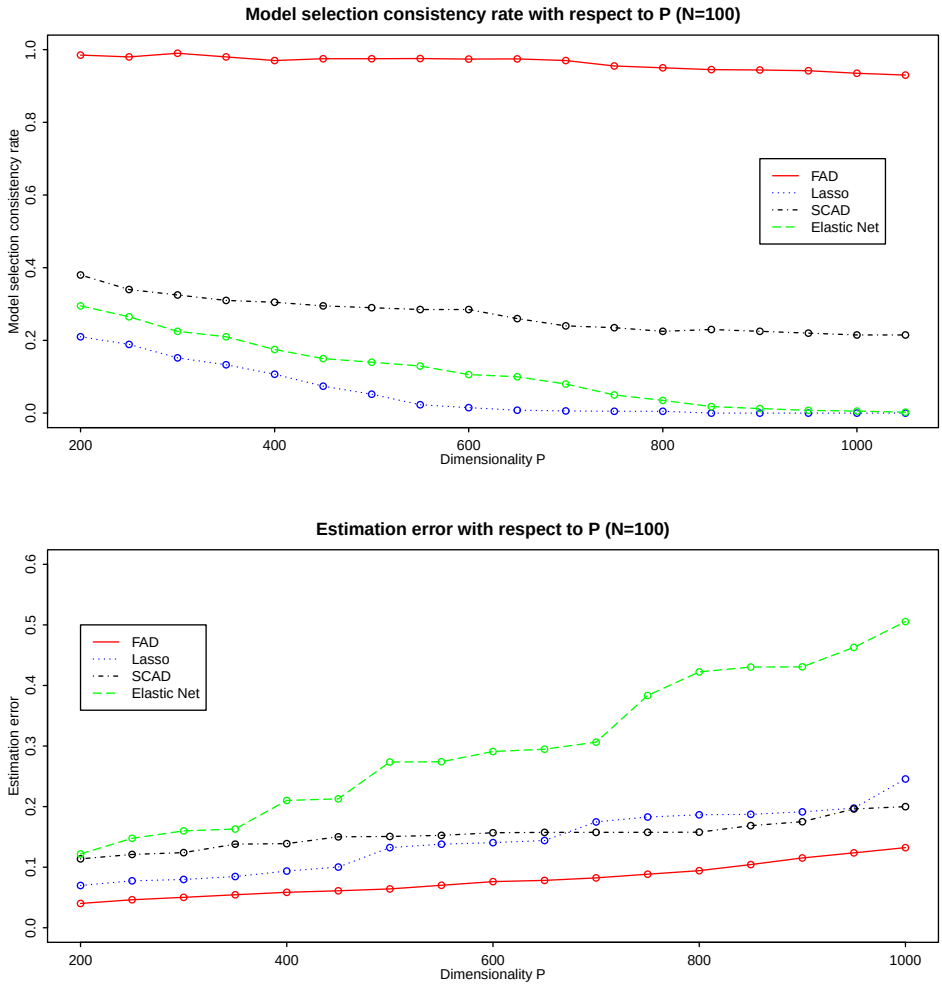

Figure 4: From above to blow: (a) Model selection consistency rate with fixed n and increasing p; (b) Mean estimation error $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\|_2$ with fixed n and increasing p.

(2) Equal correlated case. Let $\mathbf{X} = (\mathbf{x}_1, \cdots, \mathbf{x}_p) \sim N_p(\mathbf{0}, \boldsymbol{\Sigma})$, where $\boldsymbol{\Sigma}$ has diagonal elements 1 and off-diagonal elements 0.8.

(3) Uncorrelated case. Let $\mathbf{X} = (\mathbf{x}_1, \cdots, \mathbf{x}_p) \sim N_p(\mathbf{0}, \mathbf{I})$.

We compare the model selection performance of FAD with Lasso and simulate 100 replications for each scenario. The model selection performance is measured by selection consistency rate, sure screening rate and the average size of selected model. The results are presented in Table 2 below. According to Table 2, FAD pays no price for the uncorrelated case and outperforms Lasso for highly correlated cases.

Table 2: Model selection results of logistic regression ($n = 200$)

	Factor model with $K = 3$					
	FAD			Lasso		
	Selection rate	Screening rate	Average model size	Selection rate	Screening rate	Average model size
$p = 300$	0.96	1.00	3.04	0.09	0.99	12.36
$p = 400$	0.91	0.99	3.09	0.07	0.98	12.66
$p = 500$	0.86	0.98	3.13	0.06	0.96	14.83
	Equal correlated case					
	FAD			Lasso		
	Selection rate	Screening rate	Average model size	Selection rate	Screening rate	Average model size
$p = 300$	0.93	1.00	3.09	0.07	0.85	9.90
$p = 400$	0.89	1.00	3.14	0.05	0.80	10.82
$p = 500$	0.85	0.99	3.19	0.02	0.69	11.79
	Uncorrelated case					
	FAD			Lasso		
	Selection rate	Screening rate	Average model size	Selection rate	Screening rate	Average model size
$p = 300$	0.97	1.00	3.03	0.95	1.00	3.14
$p = 400$	0.93	1.00	3.07	0.91	1.00	3.34
$p = 500$	0.91	1.00	3.10	0.89	1.00	3.42

6 Real data analysis

Oberthuer et al. (2006) analyzed the German Neuroblastoma Trials NB90-NB2004 (diagnosed between 1989 and 2004) and developed a gene expression based classifier. For 246 neuroblastoma patients, gene expressions over 10,707 probe sites were measured. The binary response variable is the 3-year event-free survival information of the patients (56 positive and 190 negative). We refer to Oberthuer et al. (2006) for more biological background of the study.

We model the dataset by the high dimensional sparse logistic regression model and apply FAD to select a model. FAD selects a model with 17 covariates. In contrast, Lasso, SCAD and elastic net (with $\lambda_1 = \lambda_2 \equiv \lambda$) select much larger models including 40, 34 and 86 variables respectively. These methods may over-fit the model as they ignore the correlations among covariates. Indeed, the covariates in the dataset are highly correlated and can be approximated well by a factor model. To see this, we apply principal component analysis to the covariates and report the scree plot of the top 20 principal components in Figure 5. The scree plot shows the top ten principal components together can explain more than 50% of the total variance. The ratio method (Lam and Yao, 2012) suggests a $\widehat{K} = 4$ factor model.

To assess the performance of each model selection method, we apply a bootstrap based out-of-sample prediction as follows. For each replication, we randomly draw 200 observations as the training

set and leave the remaining 46 observations as the testing set. We use the training set to select and fit a model. Then, for each observation in the testing set, we use the fitted model to calculate the conditional probabilities that this observation belongs to different risk groups (positive or negative). The observation will be classified into the group with higher estimated conditional probability. We record the selected model size and the correct prediction rate (# of correct predictions/46). By repeating this procedure over 2,000 replications, we report the average model sizes and average correct prediction rates in Table 3. According to Table 3, FAD has the smallest average sample size and the highest correct prediction rate. Lasso, SCAD and elastic net tend to select much larger models and result in lower correct prediction rates. In Table 3 we also report the out-of-sample correct prediction rate using the first 17 variables that entered the solution path of each model selection method. As a result, the correct prediction rates of Lasso, SCAD and elastic net drop. This indicates the Lasso, SCAD and elastic net are very likely to select overfitted models.

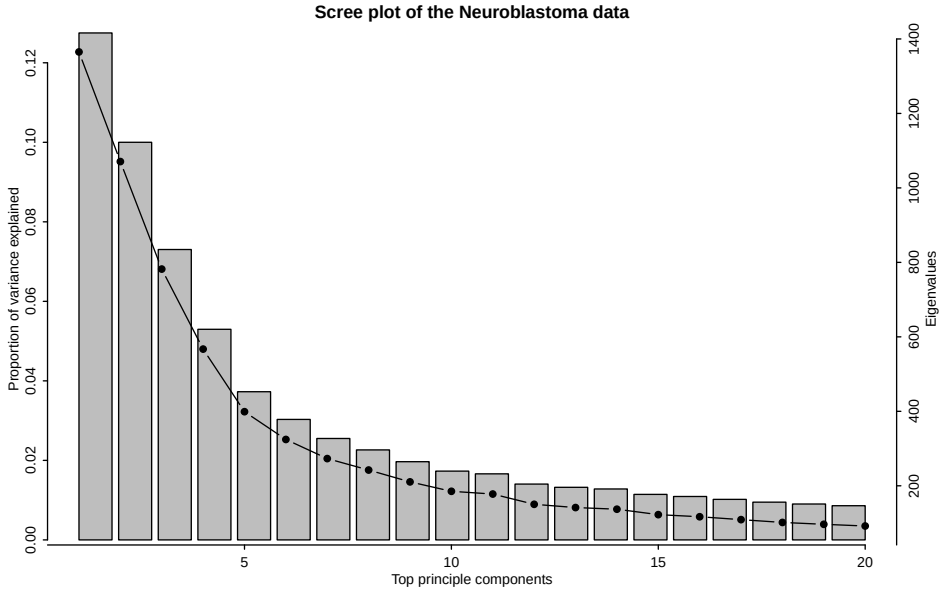

Figure 5: Eigenvalues (dotted line) and proportion of variance explained (bar) by the top 20 principal components

Table 3: Bootstrapping model selection performance of Neuroblastoma data

Bootstrap sample average	Model selection methods			
	FAD	Lasso	SCAD	elastic net
Model size	**17.6**	46.2	34.0	90.0
Correct prediction rate	**0.813**	0.807	0.809	0.790
Prediction performance with first 17 variables enter the solution path				
	FAD	Lasso	SCAD	elastic net
Correct prediction rate	**0.813**	0.733	0.764	0.705

A Some preliminary results

In the first appendix, we introduce some useful results in convex analysis and inverse problems. Under mild conditions, the tools we developed connect the unique global optimum of the regularized loss function $L_\lambda(\boldsymbol{\theta}) = L(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})$ with the solution of a constrained problem $\min_{\text{supp}(\boldsymbol{\theta})\subseteq \text{S}} L_\lambda(\boldsymbol{\theta})$.

Lemma A.1. *Suppose $L(\boldsymbol{\theta}) \in C^2(\mathbb{R}^p)$ and is convex. $R(\boldsymbol{\theta})$ is convex and $R(\boldsymbol{\alpha}+\boldsymbol{\beta}) = R(\boldsymbol{\alpha}) + R(\boldsymbol{\beta})$ for $\boldsymbol{\alpha} \in \mathcal{M}$ and $\boldsymbol{\beta} \in \mathcal{M}^\perp$, where $\mathcal{M}$ is a linear subspace of $\mathbb{R}^p$ and $\mathcal{M}^\perp$ is its orthonormal complement. Besides, there exists $R^*(\boldsymbol{\theta}) \in C(\mathbb{R}^p)$ such that $|\langle \boldsymbol{\alpha}, \boldsymbol{\beta}\rangle| \le R(\boldsymbol{\alpha})R^*(\boldsymbol{\beta})$ for $\boldsymbol{\alpha} \in \mathcal{M}^\perp$ and $\boldsymbol{\beta} \in \mathbb{R}^p$. Let $L_\lambda(\boldsymbol{\theta}) = L(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})$ where $\lambda \ge 0$, and $\widehat{\boldsymbol{\theta}} \in \text{argmin}_{\boldsymbol{\theta}\in\mathcal{M}} L_\lambda(\boldsymbol{\theta})$.*

If $R^*(\nabla L(\widehat{\boldsymbol{\theta}})) < \lambda$ and $\boldsymbol{\theta}^T\nabla^2 L(\widehat{\boldsymbol{\theta}})\boldsymbol{\theta} > 0$ for all $\boldsymbol{\theta} \in \mathcal{M}$, then $\widehat{\boldsymbol{\theta}}$ is the unique global minimizer of $L_\lambda(\boldsymbol{\theta})$.

Proof. For any $\boldsymbol{\theta} \in \mathbb{R}^p$ we use $\boldsymbol{\theta}_\mathcal{M}$, $\boldsymbol{\theta}_{\mathcal{M}^\perp}$ to denote its orthonormal projections on $\mathcal{M}$ and $\mathcal{M}^\perp$, respectively. On the one hand, by convexity and orthogonality we have

$$L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_\mathcal{M}) \ge \langle \nabla L(\boldsymbol{\theta}_\mathcal{M}), \boldsymbol{\theta} - \boldsymbol{\theta}_\mathcal{M}\rangle = \langle \nabla L(\boldsymbol{\theta}_\mathcal{M}), \boldsymbol{\theta}_{\mathcal{M}^\perp}\rangle \ge -R(\boldsymbol{\theta}_{\mathcal{M}^\perp})R^*(\nabla L(\boldsymbol{\theta}_\mathcal{M})).$$

Since $R^*(\nabla L(\widehat{\boldsymbol{\theta}})) < \lambda$, there exists $\delta > 0$ such that $\|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2 < \delta$ implies $R^*(\nabla L(\boldsymbol{\theta})) < \lambda$. Together with $\|\boldsymbol{\theta}_\mathcal{M} - \widehat{\boldsymbol{\theta}}\|_2 \le \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2$, we know $L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_\mathcal{M}) \ge -\lambda R(\boldsymbol{\theta}_{\mathcal{M}^\perp})$ as long as $\|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2 < \delta$, and the inequality strictly holds when $R(\boldsymbol{\theta}_{\mathcal{M}^\perp}) > 0$.

On the other hand, $R(\boldsymbol{\theta}) - R(\boldsymbol{\theta}_\mathcal{M}) = R(\boldsymbol{\theta}_\mathcal{M} + \boldsymbol{\theta}_{\mathcal{M}^\perp}) - R(\boldsymbol{\theta}_\mathcal{M}) = R(\boldsymbol{\theta}_{\mathcal{M}^\perp})$. Hence $\|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2 < \delta$ forces $L_\lambda(\boldsymbol{\theta}) - L_\lambda(\boldsymbol{\theta}_\mathcal{M}) = [L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_\mathcal{M})] + \lambda[R(\boldsymbol{\theta}) - R(\boldsymbol{\theta}_\mathcal{M})] \ge 0$ and the inequality strictly holds when $R(\boldsymbol{\theta}_{\mathcal{M}^\perp}) > 0$.

Now suppose $0 < \|\boldsymbol{\theta} - \widehat{\boldsymbol{\theta}}\|_2 < \delta$. If $\boldsymbol{\theta} \in \mathcal{M}$, then the facts $\widehat{\boldsymbol{\theta}} \in \text{argmin}_{\boldsymbol{\theta}\in\mathcal{M}} L_\lambda(\boldsymbol{\theta})$ and $\boldsymbol{\alpha}^T\nabla^2 L(\widehat{\boldsymbol{\theta}})\boldsymbol{\alpha} > 0$, $\forall \boldsymbol{\alpha} \in \mathcal{M}$ implies that $L_\lambda(\boldsymbol{\theta}') > L_\lambda(\widehat{\boldsymbol{\theta}})$. Besides, our assumptions yield $\|\boldsymbol{\theta}\|_2^2 \le R(\boldsymbol{\theta})R^*(\boldsymbol{\theta})$ for $\boldsymbol{\theta} \in \mathcal{M}^\perp$, leading to $R(\boldsymbol{\theta}) > 0$ over $\mathcal{M}^\perp\backslash\{0\}$. If $\boldsymbol{\theta} \notin \mathcal{M}$, then $R(\boldsymbol{\theta}_{\mathcal{M}^\perp}) > 0$ and $L_\lambda(\boldsymbol{\theta}) > L_\lambda(\boldsymbol{\theta}_\mathcal{M}) \ge L_\lambda(\widehat{\boldsymbol{\theta}})$. Therefore $\widehat{\boldsymbol{\theta}}$ is a strict local optimum of $L_\lambda(\boldsymbol{\theta})$, which is convex over $\mathbb{R}^p$. This finishes the proof.

□

Lemma A.2. *Let $L(\boldsymbol{\theta})$ be convex over a Euclidean space $\mathcal{M}$. If $\boldsymbol{\theta}_0 \in \mathcal{M}$, $r > 0$, and $L(\boldsymbol{\theta}) > L(\boldsymbol{\theta}_0)$ over the sphere $\partial B(\boldsymbol{\theta}_0, r)$, then any minimizer of $L(\boldsymbol{\theta})$ is within the ball $B(\boldsymbol{\theta}_0, r)$.*

Proof. For any $\boldsymbol{\theta} \notin \overline{B(\boldsymbol{\theta}_0, r)}$, there exists $t \in (0,1)$ and $\boldsymbol{\theta}' \in \partial B(\boldsymbol{\theta}_0, r)$ such that $\boldsymbol{\theta}' = (1-t)\boldsymbol{\theta} + t\boldsymbol{\theta}_0$. Then $L(\boldsymbol{\theta}_0) < L(\boldsymbol{\theta}') \leq (1-t)L(\boldsymbol{\theta}) + tL(\boldsymbol{\theta}_0)$, yielding $L(\boldsymbol{\theta}) > L(\boldsymbol{\theta}_0)$. Hence there is no minimizer outside $B(\boldsymbol{\theta}_0, r)$. □

Lemma A.3. *Suppose $\mathcal{M}$ is a Euclidean space, $\boldsymbol{\theta}_0 \in \mathcal{M}$ and $L(\boldsymbol{\theta})$ is convex over $\mathcal{M}$. Besides, there exist $\kappa, A > 0$ such that*

$$L(\boldsymbol{\theta}) \geq L(\boldsymbol{\theta}_0) + \langle \mathbf{h}, \boldsymbol{\theta} - \boldsymbol{\theta}_0 \rangle + \frac{\kappa}{2}\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2^2$$

as long as $\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)$ and $\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq A$. If $\inf_{\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)} \|\mathbf{h}\|_2 < \frac{1}{2}\kappa A$, then any minimizer of $L_\lambda(\boldsymbol{\theta}) = L(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})$ is within the ball $\{\boldsymbol{\theta} : \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq \frac{2}{\kappa} \inf_{\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)} \|\mathbf{h}\|_2\}$.

Proof. If $\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 < A$ and $\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)$, then

$$\begin{aligned} &L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_0) \geq \langle \mathbf{h}, \boldsymbol{\theta} - \boldsymbol{\theta}_0 \rangle + \frac{\kappa}{2}\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2^2 \geq -\|\mathbf{h}\|_2\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 + \frac{\kappa}{2}\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2^2 \\ &= \|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 (\frac{\kappa}{2}\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 - \|\mathbf{h}\|_2). \end{aligned}$$

Taking $\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)$ and $r > 0$ such that $\frac{2}{\kappa}\|\mathbf{h}\|_2 < r < A$. This forces $L(\boldsymbol{\theta}) - L(\boldsymbol{\theta}_0) > 0$ over the sphere $B(\boldsymbol{\theta}_0, r)$. Let $\widehat{\boldsymbol{\theta}}$ be one of the minimizers of $L(\boldsymbol{\theta})$. Lemma A.2 implies that $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\| < r < A$. Then

$$0 \geq L(\widehat{\boldsymbol{\theta}}) - L(\boldsymbol{\theta}_0) \geq \|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 (\frac{\kappa}{2}\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 - \|\mathbf{h}\|_2). \tag{A.1}$$

The result is proved by taking the infimum over $\mathbf{h} \in \partial L(\boldsymbol{\theta}_0)$. □

Corollary A.1. *Suppose $\lambda \geq 0$, $\mathcal{M}$ is a Euclidean space, $\boldsymbol{\theta}_0 \in \mathcal{M}$, $L(\boldsymbol{\theta}) \in C^2(\mathcal{M})$ and is convex, and $R(\boldsymbol{\theta})$ is convex. Besides, there exist $\kappa, A > 0$ such that $\nabla^2 L(\boldsymbol{\theta}) \succeq \kappa I$ as long as $\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2 \leq A$. If $\|\nabla L(\boldsymbol{\theta}_0)\|_2 + \lambda \inf_{\mathbf{h} \in \partial R(\boldsymbol{\theta}_0)} \|\mathbf{h}\|_2 < \frac{1}{2}\kappa A$, then $L_\lambda(\boldsymbol{\theta}) = L(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})$ has unique minimizer $\widehat{\boldsymbol{\theta}}$ and $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 \leq \frac{2}{\kappa}(\|\nabla L(\boldsymbol{\theta}_0)\|_2 + \lambda \inf_{\mathbf{h} \in \partial R(\boldsymbol{\theta}_0)} \|\mathbf{h}\|_2)$.*

Proof. Note that $\partial L_\lambda(\boldsymbol{\theta}_0) = \nabla L(\boldsymbol{\theta}_0) + \lambda \partial R(\boldsymbol{\theta}_0)$. There exists $\mathbf{h} \in \partial R(\boldsymbol{\theta}_0)$ such that $\mathbf{h}' = \nabla L(\boldsymbol{\theta}_0) + \mathbf{h} \in \partial L_\lambda(\boldsymbol{\theta}_0)$ and $\|\mathbf{h}'\|_2 < \|\nabla L(\boldsymbol{\theta}_0)\|_2 + \lambda\|\mathbf{h}\|_2 < \frac{1}{2}\kappa A$. Applying Lemma A.3 to L_λ and $\mathbf{h}'$, we obtain that any minimizer of L_λ satisfies $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 \leq \frac{2}{\kappa}\|\mathbf{h}'\|_2 < \frac{2}{\kappa}(\|\nabla L(\boldsymbol{\theta}_0)\|_2 + \lambda\|\mathbf{h}\|_2)$. Then $\|\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}_0\|_2 \leq A$ and $\nabla^2 L(\widehat{\boldsymbol{\theta}}) \succ 0$, proving both the bound and uniqueness. □

B Proofs of Section 4

B.1 Proof of Theorem 4.1

Define $B_S(\boldsymbol{\theta}^*, r) = \{\boldsymbol{\theta} : \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\| \leq r, \text{supp}(\boldsymbol{\theta}) \subseteq S\}$ for $r > 0$. First we introduce two useful lemmas.

Lemma B.1. *Suppose* $\mathbf{A} \in \mathbb{R}^{q\times r}$ *and* $\mathbf{B}, \mathbf{C} \in \mathbb{R}^{r\times r}$ *and* $\|\mathbf{C}\mathbf{B}^{-1}\| < 1$, *where* $\|\cdot\|$ *is an induced norm. Then* $\|\mathbf{A}[(\mathbf{B}+\mathbf{C})^{-1} - \mathbf{B}^{-1}]\| \leq \frac{\|\mathbf{A}\mathbf{B}^{-1}\|\cdot\|\mathbf{C}\mathbf{B}^{-1}\|}{1-\|\mathbf{C}\mathbf{B}^{-1}\|}$.

Proof. By the sub-multiplicity of induced norms,

$$\begin{aligned}
&\|\mathbf{A}[(\mathbf{B}+\mathbf{C})^{-1} - \mathbf{B}^{-1}]\| = \|\mathbf{A}\mathbf{B}^{-1}[(\mathbf{I}+\mathbf{C}\mathbf{B}^{-1})^{-1} - \mathbf{I}]\| \leq \|\mathbf{A}\mathbf{B}^{-1}\| \cdot \|(\mathbf{I}+\mathbf{C}\mathbf{B}^{-1})^{-1} - \mathbf{I}\| \\
&= \|\mathbf{A}\mathbf{B}^{-1}\| \cdot \Big\| \sum_{k=0}^{\infty} (-\mathbf{C}\mathbf{B}^{-1})^k - \mathbf{I} \Big\| \leq \|\mathbf{A}\mathbf{B}^{-1}\| \sum_{k=1}^{\infty} \|\mathbf{C}\mathbf{B}^{-1}\|^k = \frac{\|\mathbf{A}\mathbf{B}^{-1}\| \cdot \|\mathbf{C}\mathbf{B}^{-1}\|}{1-\|\mathbf{C}\mathbf{B}^{-1}\|}.
\end{aligned}$$

□

Lemma B.2. *Under Assumptions 4.1 and 4.2, we have* $\|(\nabla^2_{SS} L_n(\boldsymbol{\theta}))^{-1}\|_2 < \kappa_2^{-1}$ *and* $\|(\nabla^2_{SS} L_n(\boldsymbol{\theta}))^{-1}\|_\infty < \kappa_\infty^{-1}$ *over* $B_S(\boldsymbol{\theta}^*, \min\{A, \frac{\kappa_\infty}{M}\})$.

Proof. Define $\alpha_p(\boldsymbol{\theta}) = \|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}[\nabla^2_{SS} L_n(\boldsymbol{\theta}) - \nabla^2_{SS} L_n(\boldsymbol{\theta}^*)]\|_p$ for $p \in \{2, \infty\}$ and $\boldsymbol{\theta} \in B_S(\boldsymbol{\theta}^*, A)$. Note that for any symmetric matrix $\mathbf{A}$, we have $\|\mathbf{A}\|_1 = \|\mathbf{A}\|_\infty$ and $\|\mathbf{A}\|_2 \leq \sqrt{\|\mathbf{A}\|_1 \|\mathbf{A}\|_\infty} \leq \|\mathbf{A}\|_\infty$. Hence by the Assumptions we obtain that when $\|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \leq \min\{A, \frac{\kappa_\infty}{M}\}$ and $p \in \{2, \infty\}$,

$$\alpha_p(\boldsymbol{\theta}) \leq \|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_\infty \|\nabla^2_{SS} L_n(\boldsymbol{\theta}) - \nabla^2_{SS} L_n(\boldsymbol{\theta}^*)\|_\infty \leq \frac{1}{2\kappa_\infty} M \|\boldsymbol{\theta} - \boldsymbol{\theta}^*\|_2 \leq \frac{1}{2}.$$

Lemma B.1 leads to

$$\begin{aligned}
\|(\nabla^2_{SS} L_n(\boldsymbol{\theta}))^{-1} - (\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_\infty &\leq \|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_\infty \frac{\alpha_\infty}{1-\alpha_\infty} < \frac{1}{2\kappa_\infty}, \\
\|(\nabla^2_{SS} L_n(\boldsymbol{\theta}))^{-1} - (\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_2 &\leq \|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_2 \frac{\alpha_2}{1-\alpha_2} < \frac{1}{2\kappa_2}.
\end{aligned}$$

Then the proof is finished by triangle's inequality and Assumption 4.2. □

Now we are ready to prove Theorem 4.1.

Proof of Theorem 4.1. First we study the minimizer of restricted problem

$$\bar{\boldsymbol{\theta}} = \underset{\boldsymbol{\theta}\in\mathcal{M}}{\text{argmin}}\{L_n(\boldsymbol{\theta}) + \lambda R(\boldsymbol{\theta})\}.$$

Take $R(\boldsymbol{\theta}) = \|\boldsymbol{\theta}_{S_{[K]^c}}\|_1$ and $R^*(\boldsymbol{\theta}) = \|\boldsymbol{\theta}_{S_2}\|_\infty$. Let $A_1 = \min\{A, \frac{\kappa_\infty \tau}{3M}\}$ and hence $A_1 \leq \min\{A, \frac{\kappa_\infty}{M}\}$. Lemma B.2 shows that $\|(\nabla^2_{SS} L_n(\boldsymbol{\theta}))^{-1}\|_2 < \kappa_2^{-1}$ and $\|(\nabla^2_{SS} L_n(\boldsymbol{\theta}))^{-1}\|_\infty < \kappa_\infty^{-1}$ over $B_S(\boldsymbol{\theta}^*, A_1)$.

Since $\mathrm{supp}(\boldsymbol{\theta}^*) \subseteq S$, any $\mathbf{h} \in \partial R(\boldsymbol{\theta}^*)$ satisfies $\|\mathbf{h}\|_2 \leq \sqrt{|S_1|}$. Therefore

$$\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_2 + \lambda \|\mathbf{h}\|_2 \leq \frac{1}{2}\kappa_2 A_1 \leq \frac{1}{2}\kappa_2 A.$$

Then Corollary A.1 implies that $\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \leq \frac{2}{\kappa_2}(\|\nabla_S L(\boldsymbol{\theta}^*)\|_2 + \lambda\sqrt{|S_1|}) \leq A_1$.

Second, we study the L^∞ bound. On the one hand, the optimality condition yields $\nabla_S L_n(\bar{\boldsymbol{\theta}}) \in \lambda \partial \|\bar{\boldsymbol{\theta}}_{[K]^c}\|_\infty$ and hence $\|\nabla_S L_n(\bar{\boldsymbol{\theta}})\|_\infty \leq \lambda$. On the other hand, by letting $\boldsymbol{\theta}_t = (1-t)\boldsymbol{\theta}^* + t\bar{\boldsymbol{\theta}}$ $(0 \leq t \leq 1)$ we have

$$\begin{aligned}
&\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*) = \int_0^1 \nabla^2_{SS} L_n(\boldsymbol{\theta}_t)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\mathrm{d}t \\
&= \nabla^2_{SS} L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) + \int_0^1 [\nabla^2_{SS} L_n(\bar{\boldsymbol{\theta}}_t) - \nabla^2_{SS} L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\mathrm{d}t.
\end{aligned}$$

Hence

$$\begin{aligned}
&\|(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) - (\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}[\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)]\|_\infty \\
&\leq \int_0^1 \|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}[\nabla^2_{SS} L_n(\bar{\boldsymbol{\theta}}_t) - \nabla^2_{SS} L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty \mathrm{d}t \\
&\leq \|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_\infty \sup_{t\in[0,1]} \|\nabla^2_{SS} L_n(\bar{\boldsymbol{\theta}}_t) - \nabla^2_{SS} L_n(\boldsymbol{\theta}^*)\|_\infty \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty
\end{aligned}$$

By Assumptions 4.1 and 4.2, we obtain that

$$\|(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) - (\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}[\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)]\|_\infty \leq \frac{M}{2\kappa_\infty}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty.$$

By $\bar{\boldsymbol{\theta}} \in B_S(\boldsymbol{\theta}^*, A_1)$ we have

$$\begin{aligned}
&\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \leq \|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_\infty \|\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \frac{M}{2\kappa_\infty}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \\
&\leq \frac{1}{2\kappa_\infty}(\lambda + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty) + \frac{1}{6}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty.
\end{aligned}$$

Therefore,

$$|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \leq \frac{3}{5\kappa_\infty}(\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \lambda). \tag{B.1}$$

Third we study the L^1 bound. The bound on $\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1$ can be obtained in a similar way. Using

the fact that $\|\cdot\|_1 = \|\cdot\|_\infty$ for symmetric matrices,

$$\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \le \|(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}\|_1 \|\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)\|_1 + \frac{M}{2\kappa_\infty}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1$$
$$\le \frac{1}{2\kappa_\infty}(\lambda|S_1| + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_1) + \frac{1}{6}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1.$$

Hence $\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \le \frac{3}{5\kappa_\infty}(\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_1 + \lambda|S_1|)$. Since $\mathrm{supp}(\bar{\boldsymbol{\theta}}) \subseteq S$, we also have

$$\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_1 \le \sqrt{|S|}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2 \le \frac{2\sqrt{|S|}}{\kappa_2}(\|\nabla_S L(\boldsymbol{\theta}^*)\|_2 + \lambda\sqrt{|S_1|}).$$

This gives another L^1 bound.

By Lemma A.1, to derive $\widehat{\boldsymbol{\theta}} = \bar{\boldsymbol{\theta}}$ it remains to show that $\|\nabla_{S_2} L_n(\bar{\boldsymbol{\theta}})\|_\infty < \lambda$. Using the Taylor expansion we have

$$\begin{aligned}\nabla_{S_2} L_n(\bar{\boldsymbol{\theta}}) - \nabla_{S_2} L_n(\boldsymbol{\theta}^*) &= \int_0^1 \nabla^2_{S_2 S} L_n(\boldsymbol{\theta}_t)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\mathrm{d}t \\ &= \nabla^2_{S_2 S} L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*) + \int_0^1 [\nabla^2_{S_2 S} L_n(\boldsymbol{\theta}_t) - \nabla^2_{S_2 S} L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\mathrm{d}t.\end{aligned} \tag{B.2}$$

On the one hand, the first term in (B.2) follows,

$$\|\nabla^2_{S_2 S} L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty = \|[\nabla^2_{S_2 S} L_n(\boldsymbol{\theta}^*)(\nabla^2_{SS} L_n(\boldsymbol{\theta}^*))^{-1}][\nabla^2_{SS} L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)]\|_\infty$$
$$\le (1-\tau)\|\nabla^2_{SS} L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty.$$

By the Taylor expansion, triangle's inequality, Assumption 4.1 and the fact that $\bar{\boldsymbol{\theta}} \in B_S(\boldsymbol{\theta}^*, A_1)$,

$$\|\nabla^2_{SS} L_n(\boldsymbol{\theta}^*)(\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty \le \|\nabla_S L_n(\bar{\boldsymbol{\theta}}) - \nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \int_0^1 \|[\nabla^2_{SS} L_n(\bar{\boldsymbol{\theta}}_t) - \nabla^2_{SS} L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\|_\infty \mathrm{d}t$$
$$\le \|\nabla_S L_n(\bar{\boldsymbol{\theta}})\|_\infty + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + M\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_2\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty$$
$$\le \lambda + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \frac{\kappa_\infty \tau}{3}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty.$$

On the other hand, we bound the second term in (B.2). Note that $\boldsymbol{\theta}_t \in B_S(\boldsymbol{\theta}^*, A_1)$ for all $t \in [0,1]$. Assumption 4.1 yields

$$\Big\| \int_0^1 [\nabla^2_{S_2 S} L_n(\boldsymbol{\theta}_t) - \nabla^2_{S_2 S} L_n(\boldsymbol{\theta}^*)](\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*)\mathrm{d}t \Big\|_\infty$$
$$\le \sup_{t\in[0,1]} \|\nabla^2_{S_2 S} L_n(\boldsymbol{\theta}_t) - \nabla^2_{S_2 S} L_n(\boldsymbol{\theta}^*)\|_\infty \|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty \le \frac{\kappa_\infty \tau}{3}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty.$$

As a result,

$$\|\nabla_{S_2} L_n(\bar{\boldsymbol{\theta}})\|_\infty \le \|\nabla_{S_2} L_n(\boldsymbol{\theta}^*)\|_\infty + (1-\tau)\Big(\lambda + \|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \frac{\kappa_\infty \tau}{3}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty\Big) + \frac{\kappa_\infty \tau}{3}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty$$
$$\le \lambda - \tau\Big(\lambda - \frac{2\kappa_\infty}{3}\|\bar{\boldsymbol{\theta}} - \boldsymbol{\theta}^*\|_\infty - \frac{2}{\tau}\|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty\Big).$$

Recall that the L^∞ bound in (B.1). By plugging in this estimate, and using the assumptions $0 < \tau < 1$ and $\lambda > \frac{20}{3\tau}\|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty$, we derive that

$$\begin{aligned}\|\nabla_{S_2} L_n(\bar{\boldsymbol{\theta}})\|_\infty &\le \lambda - \tau\Big(\lambda - \frac{2}{5}(\|\nabla_S L_n(\boldsymbol{\theta}^*)\|_\infty + \lambda) - \frac{2}{\tau}\|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty\Big)\\ &\le \lambda - \tau\Big(\frac{3}{5}\lambda - \frac{4}{\tau}\|\nabla L_n(\boldsymbol{\theta}^*)\|_\infty\Big) < \lambda.\end{aligned}$$

This implies $\widehat{\boldsymbol{\theta}} = \bar{\boldsymbol{\theta}}$ and translates all the bounds for $\bar{\boldsymbol{\theta}}$ to the ones for $\widehat{\boldsymbol{\theta}}$. The proposition on sign consistency follows from elementary computation, thus we omit its proof. □

B.2 Proof of Theorem 4.2

Proof of Theorem 4.2. Recall that $\widehat{\boldsymbol{\theta}} = \operatorname{argmin}_{\boldsymbol{\theta}}\{L_n(\mathbf{Y}, \widehat{\mathbf{W}}\boldsymbol{\theta}) + \lambda\|\boldsymbol{\theta}_{[K]^c}\|_1\}$. Also, Assumption 4.6 tells us $\mathbf{H}_0$ is nonsingular and $\mathbf{H} = \begin{pmatrix} \mathbf{H}_0 & \mathbf{0}_{K\times p} \\ \mathbf{0}_{p\times K} & \mathbf{I}_p \end{pmatrix}$. Define $\overline{\mathbf{W}} = \widehat{\mathbf{W}}\mathbf{H}$, $\bar{\boldsymbol{\theta}} = \mathbf{H}^{-1}\widehat{\boldsymbol{\theta}}$ and $\bar{\boldsymbol{\theta}}^* = \mathbf{H}^{-1}\boldsymbol{\theta}^*$. We easily see that $\widehat{\boldsymbol{\beta}} = \widehat{\boldsymbol{\theta}}_{[K]^c} = \bar{\boldsymbol{\theta}}_{[K]^c}$ and $\bar{\boldsymbol{\theta}} = \operatorname{argmin}_{\boldsymbol{\theta}}\{L_n(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta}) + \lambda\|\boldsymbol{\theta}_{[K]^c}\|_1\}$. Then it follows that $\operatorname{supp}(\widehat{\boldsymbol{\beta}}) = \operatorname{supp}(\bar{\boldsymbol{\theta}}_{[K]^c})$ and $\|\widehat{\boldsymbol{\beta}} - \boldsymbol{\beta}^*\| = \|\bar{\boldsymbol{\theta}}_{[K]^c} - \bar{\boldsymbol{\theta}}^*_{[K]^c}\| \le \|\bar{\boldsymbol{\theta}} - \bar{\boldsymbol{\theta}}^*\|$ for any norm $\|\cdot\|$.

Consequently, Theorem 4.2 is reduced to studying $\bar{\boldsymbol{\theta}}$ and the loss function $L_n(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta})$. The Lemma B.3 below implies that all the regularity conditions (with $A = \infty$) in Theorem 4.1 are satisfied.

We claim that $\|\nabla L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_\infty \le \varepsilon$, and below is the proof. Let w_{ij} and $\overline{w}_{ij}$ be the (i,j)-th element of $\mathbf{W}$ and $\overline{\mathbf{W}}$, respectively. Note that $L_n(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta}) = \frac{1}{n}\sum_{i=1}^n[-y_i\overline{\mathbf{w}}_i^T\boldsymbol{\theta} + b(\overline{\mathbf{w}}_i^T\boldsymbol{\theta})]$. The chain rule of differentiation yields

$$\nabla L_n(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta}) = \frac{1}{n}\sum_{i=1}^n[-y_i + b'(\overline{\mathbf{w}}_i^T\boldsymbol{\theta})]\overline{\mathbf{w}}_i.$$

Note that $\overline{\mathbf{w}}_i^T\bar{\boldsymbol{\theta}}^* = \mathbf{x}_i^T\boldsymbol{\beta}^*$. For $j \in [K+p]$, we have

$$\frac{\partial L_n}{\partial \theta_j}(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) = \frac{1}{n}\sum_{i=1}^n \overline{w}_{ij}[-y_i + b'(\mathbf{x}_i^T\boldsymbol{\beta}^*)].$$

Hence $\|\nabla L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_\infty \le \varepsilon$ and consequently, $\|\nabla_S L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_\infty \le \varepsilon$, $\|\nabla_S L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_2 \le$

$\varepsilon\sqrt{|S|}$ and $\|\nabla_S L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_1 \leq \varepsilon|S|$. Besides, $\lambda > 7\varepsilon/\tau \geq \varepsilon$. Based on these estimates, all the results follow from Theorem 4.1 and some simple algebra. □

Here we present the Lemma B.3 used above and its proof.

Lemma B.3. *Let Assumptions 4.4, 4.5 and 4.6 hold. Treat $L_n(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta})$ as a function of $\boldsymbol{\theta}$, and the derivatives below are taken with respect to it. Define $M = M_0^3 M_3 |S|^{3/2}$. Then*

$$
\begin{aligned}
&(i) \quad \|\nabla^2_{\cdot S} L(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta}) - \nabla^2_{\cdot S} L(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_\infty \leq M\|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2, \ \forall \boldsymbol{\theta}, \\
&(ii) \quad \|(\nabla^2_{SS} L(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \leq \frac{1}{2\kappa_\infty}, \\
&(iii) \quad \|(\nabla^2_{SS} L(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_2 \leq \frac{1}{2\kappa_2}, \\
&(iv) \quad \|\nabla^2_{S_2 S} L(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)(\nabla^2_{SS} L(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \leq 1 - \tau.
\end{aligned}
$$

Proof. (i) The chain rule of differentiation yields

$$
\begin{aligned}
\nabla^2 L_n(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta}) &= \frac{1}{n}\sum_{i=1}^n b''(\overline{\mathbf{w}}_i^T\boldsymbol{\theta})\overline{\mathbf{w}}_i\overline{\mathbf{w}}_i^T, \\
\nabla^2 L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}) &= \frac{1}{n}\sum_{i=1}^n b''(\mathbf{w}_i^T\boldsymbol{\theta})\mathbf{w}_i\mathbf{w}_i^T.
\end{aligned}
$$

Based on the fact that $\mathbf{w}_i^T\boldsymbol{\theta}^* = \overline{\mathbf{w}}_i^T\bar{\boldsymbol{\theta}}^* = \mathbf{x}_i^T\boldsymbol{\beta}^*$, we have

$$
\begin{aligned}
\nabla^2 L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*) &= \frac{1}{n}\sum_{i=1}^n b''(\mathbf{x}_i^T\boldsymbol{\beta}^*)\mathbf{w}_i\mathbf{w}_i^T, \\
\nabla^2 L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) &= \frac{1}{n}\sum_{i=1}^n b''(\mathbf{x}_i^T\boldsymbol{\beta}^*)\overline{\mathbf{w}}_i\overline{\mathbf{w}}_i^T.
\end{aligned}
$$

For any $j, k \in [p+K]$ and $\mathrm{supp}(\boldsymbol{\theta}) \subseteq S$,

$$
\begin{aligned}
&|\nabla^2_{jk} L_n(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta}) - \nabla^2_{jk} L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)| \leq \frac{1}{n}\sum_{i=1}^n |b''(\overline{\mathbf{w}}_i^T\boldsymbol{\theta}) - b''(\overline{\mathbf{w}}_i^T\bar{\boldsymbol{\theta}}^*)| \cdot |\overline{\mathbf{w}}_{ij}\overline{\mathbf{w}}_{ik}| \\
&\leq \frac{1}{n}\sum_{i=1}^n M_3 |\overline{\mathbf{w}}_i^T(\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*)| \cdot \|\overline{\mathbf{W}}\|^2_{\max}
\end{aligned} \tag{B.3}
$$

By the Cauchy-Schwarz inequality and $\|\overline{\mathbf{W}}\|_{\max} \leq \|\mathbf{W}\|_{\max} + \|\overline{\mathbf{W}} - \mathbf{W}\|_{\max} \leq M_0$, we obtain that for $i \in [n]$,

$$
|\overline{\mathbf{w}}_i^T(\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*)| = |\overline{\mathbf{w}}_{iS}^T(\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*)_S| \leq \|\overline{\mathbf{w}}_{iS}\|_2\|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2 \leq \sqrt{|S|}\|\overline{\mathbf{W}}\|_{\max}\|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2.
$$

Plugging this result back to (B.3), we get

$$|\nabla^2_{jk} L_n(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta}) - \nabla^2_{jk} L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)| \le \sqrt{|S|} M_3 M_0^3 \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2, \quad \forall j, k \in [K+p];$$
$$\|\nabla^2_{\cdot S} L_n(\mathbf{Y}, \overline{\mathbf{W}}\boldsymbol{\theta}) - \nabla^2_{\cdot S} L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)\|_\infty \le |S|^{3/2} M_3 M_0^3 \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2 = M \|\boldsymbol{\theta} - \bar{\boldsymbol{\theta}}^*\|_2.$$

(ii) Now we come to the second claim. For any $k \in [p+K]$,

$$\|\nabla^2_{kS} L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla^2_{kS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*)\|_\infty \le \frac{1}{n} \sum_{i=1}^n b''(\mathbf{x}_i^T \boldsymbol{\beta}^*) \|\overline{\mathbf{w}}_{ik} \overline{\mathbf{w}}_{iS}^T - \mathbf{w}_{ik} \mathbf{w}_{iS}^T\|_\infty$$
$$\le \frac{M_2 \sqrt{|S|}}{n} \sum_{i=1}^n \|\overline{\mathbf{w}}_{ik} \overline{\mathbf{w}}_{iS}^T - \mathbf{w}_{ik} \mathbf{w}_{iS}^T\|_2.$$

Also, by $\|\mathbf{W}\|_{\max} \le M_0/2$ and $\|\overline{\mathbf{W}}\|_{\max} \le M_0$ we have

$$\begin{aligned}
&\|\overline{\mathbf{w}}_{ik} \overline{\mathbf{w}}_{iS}^T - \mathbf{w}_{ik} \mathbf{w}_{iS}^T\|_2 \le |\mathbf{w}_{ik}| \cdot \|(\overline{\mathbf{w}}_{iS} - \mathbf{w}_{iS})^T\|_2 + |\overline{\mathbf{w}}_{ik} - \mathbf{w}_{ik}| \cdot \|\overline{\mathbf{w}}_{iS}^T\|_2 \\
&\le \|\mathbf{W}\|_{\max} \|\overline{\mathbf{w}}_{iS} - \mathbf{w}_{iS}\|_2 + |\overline{\mathbf{w}}_{ik} - \mathbf{w}_{ik}| \cdot \sqrt{|S|} \|\overline{\mathbf{W}}\|_{\max} \\
&\le \frac{M_0}{2} \|\overline{\mathbf{w}}_{iS} - \mathbf{w}_{iS}\|_2 + M_0 \sqrt{|S|} \cdot |\overline{\mathbf{w}}_{ik} - \mathbf{w}_{ik}|.
\end{aligned}$$

Define $\delta = \max_{1 \le j \le p+K} (\frac{1}{n} \sum_{i=1}^n |\overline{w}_{ij} - w_{ij}|^2)^{1/2}$. By the Jensen's inequality, $\forall J \subseteq [p+K]$,

$$\frac{1}{n} \sum_{i=1}^n \|\overline{\mathbf{w}}_{iJ} - \mathbf{w}_{iJ}\|_2 \le \Big(\frac{1}{n} \sum_{i=1}^n \|\overline{\mathbf{w}}_{iJ} - \mathbf{w}_{iJ}\|_2^2 \Big)^{1/2} \le \Big(\frac{|J|}{n} \max_{1 \le j \le K+p} \sum_{i=1}^n |\overline{w}_{ij} - w_{ij}|^2 \Big)^{1/2} \le \sqrt{|J|} \delta.$$

As a result,

$$\begin{aligned}
&\|\nabla^2_{\cdot S} L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla^2_{\cdot S} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*)\|_\infty \\
&= \max_{k \in [K+p]} \|\nabla^2_{kS} L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla^2_{kS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*)\|_\infty \le \frac{3}{2} M_0 M_2 |S| \delta.
\end{aligned}$$

Let $\alpha = \|(\nabla^2_{SS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1} [\nabla^2_{SS} L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*) - \nabla^2_{SS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*)]\|_\infty$. Then

$$\begin{aligned}
\alpha &\le \|(\nabla^2_{SS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \|\nabla^2_{SS} L_n(\mathbf{Y}, \overline{\mathbf{W}}^T \bar{\boldsymbol{\theta}}^*) - \nabla^2_{SS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*)\|_\infty \\
&\le \frac{3}{8\kappa_\infty} M_0 M_2 |S| \delta \le \frac{1}{2}.
\end{aligned} \tag{B.4}$$

Lemma B.1 yields

$$\|(\nabla^2_{SS} L_n(\mathbf{Y}, \overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - (\nabla^2_{SS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \le \|(\nabla^2_{SS} L_n(\mathbf{Y}, \mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \frac{\alpha}{1-\alpha}$$
$$\le \frac{1}{4\kappa_\infty} \cdot \frac{\alpha}{1 - \frac{1}{2}} \le \frac{3}{16\kappa_\infty^2} M_0 M_2 |S| \delta.$$

We also have a cruder bound $\|(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}-(\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty \le \frac{1}{4\kappa_\infty}$, which leads to

$$\|(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \le \|(\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty + \frac{1}{4\kappa_\infty} \le \frac{1}{2\kappa_\infty}. \tag{B.5}$$

(iii) The third argument follows (B.5) easily. Since $\|\mathbf{A}\|_2 \le \|\mathbf{A}\|_\infty$ holds for any symmetric matrix $\mathbf{A}$, we have

$$\|(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}-(\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_2 \le \frac{1}{4\kappa_\infty} \le \frac{1}{4\kappa_2}.$$

and thus $\|(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_2 \le \frac{1}{2\kappa_2}$.

(iv) Finally we prove the last inequality. On the one hand,

$$\begin{aligned}
&\|\nabla^2_{S_2S}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}-\nabla^2_{S_2S}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)(\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty\\
&\le \|\nabla^2_{S_2S}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)-\nabla^2_{S_2S}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)\|_\infty\|(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty\\
&+\|\nabla^2_{S_2S}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)[(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}-(\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*))^{-1}]\|_\infty.
\end{aligned}$$

It is easy to see that

$$\|\nabla^2_{S_2S}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)-\nabla^2_{S_2S}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)\|_\infty\|(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \le \frac{1}{4\kappa_\infty}3M_0M_2|S|\delta.$$

On the other hand, we can take $\mathbf{A}=\nabla^2_{S_2S}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)$, $\mathbf{B}=\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)$ and $\mathbf{C}=\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)-\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)$. By Assumption 4.5, $\|\mathbf{AB}^{-1}\|_\infty \le 1-2\tau \le 1$. Lemma B.1 forces that

$$\begin{aligned}
&\|\nabla^2_{S_2S}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)[(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}-(\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*))^{-1}]\|_\infty\\
&=\|\mathbf{A}[(\mathbf{B}+\mathbf{C})^{-1}-\mathbf{B}^{-1}]\|_\infty \le \|\mathbf{AB}^{-1}\|_\infty\frac{\|\mathbf{CB}^{-1}\|_\infty}{1-\|\mathbf{CB}^{-1}\|_\infty} \le \frac{\|\mathbf{C}\|_\infty\|\mathbf{B}^{-1}\|_\infty}{1-\|\mathbf{C}\|_\infty\|\mathbf{B}^{-1}\|_\infty}.
\end{aligned}$$

We have shown above in (B.4) that

$$\|\mathbf{C}\|_\infty\|\mathbf{B}^{-1}\|_\infty \le \frac{3}{8\kappa_\infty}M_0M_2|S|\delta \le \frac{1}{2}.$$

As a result,

$$\|\nabla^2_{S_2S}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)[(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}-(\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*))^{-1}]\|_\infty \le \frac{3}{4\kappa_\infty}M_0M_2|S|\delta.$$

By combining these estimates, we have

$$\|\nabla^2_{S_2S}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1} - \nabla^2_{S_2S}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*)(\nabla^2_{SS}L_n(\mathbf{Y},\mathbf{W}\boldsymbol{\theta}^*))^{-1}\|_\infty$$
$$\leq \frac{3}{2\kappa_\infty}M_0M_2|S|\delta \leq \tau.$$

Therefore $\|\nabla^2_{S_2S}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*)(\nabla^2_{SS}L_n(\mathbf{Y},\overline{\mathbf{W}}\bar{\boldsymbol{\theta}}^*))^{-1}\|_\infty \leq (1-2\tau)+\tau = 1-\tau$. □

References

Ahn, S. C. and Horenstein, A. R. (2013). Eigenvalue ratio test for the number of factors. *Econometrica*, **81**, 1203–1227.

Bai, J. and Ng, S. (2002). Determining the number of factors in approximate factor models. *Econometrica*, **70**, 191–221.

Bickel, P. J., Ritov, Y. A., and Tsybakov, A. B. (2009). Simultaneous analysis of Lasso and Dantzig selector. *The Annals of statistics*, **37**, 1705–1732.

Bunea, F. (2008). Honest variable selection in linear and logistic regression models via l_1 and l_1 + l_2 penalization. *Electronic Journal of Statistics*, **2**, 1153–1194.

Candes, E. and Tao, T. (2007). The Dantzig selector: Statistical estimation when p is much larger than n. *The Annals of Statistics*, **35**, 2313–2351.

Donoho, D. L. and Elad, M. (2003). Optimally sparse representation in general (nonorthogonal) dictionaries via l_1 minimization. *Proceedings of the National Academy of Sciences*, **100**, 2197–2202.

Efron, B., Hastie, T., Johnstone, I. and Tibshirani, R. (2004). Least angle regression. *The Annals of statistics*, **32**, 407–499.

Fan, J., and Li, R. (2001). Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*, **96**, 1348-1360.

Fan, J., Liao, Y. and Mincheva, M. (2013). Large covariance estimation by thresholding principal orthogonal complements (with discussion). *Journal of the Royal Statistical Society, Series B*, **75**, 603–680.

Fan, J. and Lv, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *Journal of the Royal Statistical Society, Series B*, **70**, 849–911.

Fan, J. and Peng, H. (2004). On non-concave penalized likelihood with diverging number of parameters. *Annals of Statistics*, **32**, 928-961.

Fan, J. and Song, R. (2009). Sure Independence Screening in Generalized Linear Models with NP-Dimensionality. *Annals of Statistics*, **38**,3567–3604.

Friedman, J., Hastie, T. and Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of statistical software*, **33**, 1–22.

Hallin, M. and Liška, R. (2007). Determining the number of factors in the general dynamic factor model. *Journal of the American Statistical Association*, **102**, 603–617.

Jia, J. and Rohe, K. (2015). Preconditioning the Lasso for sign consistency. *Electronic Journal of Statistics*, **9**, 1150–1172.

Johnstone, I. M. and Lu, A. Y. (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, **104**, 682–693.

Kneip, A. and Sarda, P. (2011). Factor models and variable selection in high-dimensional regression analysis. *The Annals of Statistics*, **39**, 2410–2447.

Lam, C. and Yao, Q. (2012). Factor modeling for high dimensional time-series: inference for the number of factors. *Annals of Statistics*, **40**, 694–726.

Lee, J. D., Sun, Y. and Taylor, J. E. (2015). On model selection consistency of regularized M-estimators. *Electronic Journal of Statistics*, **9**, 608–642.

Meinshausen, N. and Bühlmann, P. (2006). High-dimensional graphs and variable selection with the lasso. *The annals of statistics*, **34**, 1436–1462.

Negahban, S., Ravikumar, P. K., Wainwright, M. J. and Yu, B. (2012). A unified framework for high-dimensional analysis of M-estimators with decomposable regularizers. *Statistical Science*, **27**, 1348–1356.

Oberthuer, A., Berthold, F., Warnat, P., Hero, B., Kahlert, Y., Spitz, R., Ernestus, K., König, R., Haas, S.,, Eils, R., and Schwab, M. (2006). Customized oligonucleotide microarray gene expression-based classification of neuroblastoma patients outperforms current clinical risk stratification. *Journal of clinical oncology*, **24**, 5070–5078.

Paul, D., Bair, E., Hastie, T. and Tibshirani, R. (2008). "Preconditioning" for feature selection and regression in high-dimensional problems. *The Annals of Statistics*, **36**, 1595–1618.

Ravikumar, P., Wainwright, M. J., Lafferty, J. D. and Marron, J. S. (2010). High-dimensional Ising model selection using l_1-regularized logistic regression. *The Annals of Statistics*, **38**, 1287–1319.

Shen, D., Shen, H., Zhu, H. and Marron, J. S. (2016). The statistics and mathematics of high dimension low sample size asymptotics. *Statistica Sinica*, **26**, 1747–1770.

Stock, J. and Watson, M. (2002). Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association*, **97**, 1167–1179.

Su, W., Bogdan, M. and Candes, E. (2017). False discoveries occur early on the lasso paths. *The Annals of Statistics*, to appear.

Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **58**, 267–288.

Van De Geer, S. A. (2008). High-dimensional generalized linear models and the lasso. *The Annals of Statistics*, **36**, 614–645.

Van De Geer, S. A. and Müller, P. (2012). Quasi-Likelihood and/or Robust Estimation in High Dimensions *Statistical Science*, **27**, 469–480.

Wainwright, M. J. (2009). Sharp thresholds for high-dimensional and noisy sparsity recovery using-constrained quadratic programming (lasso). quadratic programming (Lasso) *IEEE transactions on information theory*, **55**, 2183–2202.

Wang, H (2012). Factor profiled sure independence screening. *Biometrika*, **99**, 15–28.

Wang, W. and Fan, J. (2017). Asymptotics of Empirical Eigen-structure for Ultra-high Dimensional Spiked Covariance Model. *Annals of Statistics*, to appear.

Wang, X. and Leng, C. (2016). High dimensional ordinary least squares projection for screening variables. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **78**, 589–611.

Zhang, C. H. (2010). Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, **38**, 894–942.

Zhao, P. and Yu, B. (2006). On model selection consistency of Lasso. *Journal of Machine Learning Research*, **7**, 2541–2563.

Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **67**, 301–320.

Zou, H. and Li, R. (2008). One-step sparse estimates in nonconcave penalized likelihood models. *The Annals of statistics*, **36**, 1509–1533.